

Dina Hagen Steinskog
Amund Fredrik Strømsnes
Lea Arwen Utstøl

An Evaluation of the Security Measures Provided by Microsoft in Azure; with Focus on Entra ID, Entra ID Protection and Sentinel: A Practical Approach

Bachelor's thesis in Digital infrastruktur og cybersikkerhet
Supervisor: Tor Ivar Melling
May 2024

Dina Hagen Steinskog
Amund Fredrik Strømsnes
Lea Arwen Utstøl

An Evaluation of the Security Measures Provided by Microsoft in Azure; with Focus on Entra ID, Entra ID Protection and Sentinel: A Practical Approach

Bachelor's thesis in Digital infrastruktur og cybersikkerhet
Supervisor: Tor Ivar Melling
May 2024

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Computer Science



An Evaluation of the Security Measures Provided by Microsoft in Azure; with Focus on Entra ID, Entra ID Protection and Sentinel: A Practical Approach

Dina Hagen Steinskog, Amund Fredrik Strømsnes, Lea Arwen Utstøl

CC-BY 2024/05/21

Abstract

Attacks targeting user accounts have long been an issue regarding online IT solutions. Azure is the most prominent cloud platform among Norwegian organisations, and Microsoft provides many security solutions and recommendations for how to secure an organisation's Azure environment.

This thesis investigates how well the security measures recommended by Microsoft, regarding Conditional Access policies and Sentinel, set up with the rule-sets provided for Entra ID and Entra ID Protection found in Sentinel Content Hub, manage to defend an organisation against attacks targeting users. We use a minimalist setup which is not tailored to a specific sector and adopt a practical approach for simulating the attacks, which are based on attack techniques observed in real-world attacks. This enables our results to be more closely related to attacks happening in the real world, and our conclusion can be used as a foundation for any organisation.

Our results show that this security solution does protect an organisation well, and also provides alerts for the organisation to be able to react to attacks in both the Sentinel and Entra ID Protection dashboards. However, the analytics rules for Sentinel do not cover all attack techniques and we recommend actively working with both dashboards if an organisation does not have the competence needed to create new rules.

Sammendrag

Angrep mot brukerekontoer har lenge vært et problem for IT-tjenester som er eksponert for internettet. I Norge er Azure den mest populære skyplattformen, og Microsoft har mange råd og anbefalinger for hvordan en organisasjon kan beskytte Azure-miljøet sitt på best mulig måte.

Målet for denne oppgaven er å utforske hvor godt rådene til Microsoft fungerer mot å beskytte en organisasjon mot angrep på brukere. Vi evaluerer hvor effektive Microsoft sine anbefalinger om Conditional Access og Sentinel er i å forsvare en organisasjon, hvor Sentinel er satt opp med regelsettene Microsoft utgir på Sentinel Content Hub for Entra ID og Entra ID Protection. Under testingen brukes det et minimalistisk Azure-oppsett og en praktisk tilnærming til simulering av angrep, som tar utgangspunkt i angrepsteknikker som har vært observert i ekte angrep tidligere. Denne tilnærmingen medfører at resultatene våre bedre reflekterer oppsettets respons på et faktisk angrep og konklusjonen vår står uavhengig av en enkelt sektor, og kan dermed benyttes av et bredere spekter av organisasjoner som et utgangspunkt for deres sikkerhetsløsning.

Våre resultater viser at denne sikkerhetsløsningen effektivt beskytter en organisasjon og gir oversikt over sikkerhetshendelser som trenger behandling i både Sentinel og Entra ID Protection sine dashbord. Likevel dekker ikke alle Sentinel-reglene alle angrepsteknikker, og vi anbefaler å jobbe aktivt med begge dashbordene dersom en organisasjon ikke sitter på kompetansen til å lage egne regler.

Contents

| | |
|-------------------------------------------------------------|-------------|
| Abstract | iii |
| Sammendrag | v |
| Contents | vii |
| Figures | xi |
| Glossary | xiii |
| Acronyms | xv |
| Preface | xvii |
| 1 Introduction | 1 |
| 1.1 Background | 1 |
| 1.2 Problem Statement | 2 |
| 1.3 Research Questions | 2 |
| 1.4 Effect Goals | 3 |
| 1.5 Project Goals | 3 |
| 1.6 Partner Organisation | 3 |
| 1.7 Thesis Outline | 4 |
| 1.8 Scope and Delimitation | 5 |
| 2 Theory | 7 |
| 2.1 Identity Access Management | 7 |
| 2.1.1 What is Identity Access Management | 7 |
| 2.1.2 Authentication and Authorisation | 8 |
| 2.1.3 How Identity Access Management Works | 8 |
| 2.1.4 Why Identity Access Management Is Important | 9 |
| 2.2 Microsoft Entra ID | 9 |
| 2.2.1 What Is Microsoft Entra ID? | 9 |
| 2.2.2 Features in Microsoft Entra ID | 10 |
| 2.2.3 Concepts in Microsoft Entra ID | 10 |
| 2.2.4 Azure Smart Lockout | 11 |
| 2.3 Microsoft Entra Identity Protection | 12 |
| 2.3.1 What Is It and What Does It Do? | 12 |
| 2.3.2 Detecting Risk | 12 |
| 2.3.3 Presenting Risk | 15 |
| 2.3.4 Simulating Risk | 15 |
| 2.4 Conditional Access | 18 |
| 2.4.1 Risk-Based Access Control Policies | 18 |

| | | |
|----------|---------------------------------------------------|-----------|
| 2.4.2 | Zero Trust | 19 |
| 2.4.3 | Best-Practice Setup | 20 |
| 2.5 | Microsoft Sentinel | 22 |
| 2.5.1 | What Is SIEM? | 22 |
| 2.5.2 | What Is SOAR? | 23 |
| 2.5.3 | SIEM and SOAR | 23 |
| 2.5.4 | How Does Microsoft Sentinel Work? | 24 |
| 2.5.5 | Content Hub | 25 |
| 2.5.6 | Sentinel Dashboard | 25 |
| 2.6 | False Results and Alert Fatigue | 26 |
| 2.7 | Attack Patterns | 26 |
| 2.7.1 | Reconnaissance | 27 |
| 2.7.2 | Initial Access | 27 |
| 2.7.3 | Credential Access | 27 |
| 2.7.4 | Privilege Escalation | 29 |
| 2.7.5 | Command and Control | 29 |
| 2.7.6 | Gaining persistence | 30 |
| 3 | Method | 31 |
| 3.1 | Outline of Chapter | 31 |
| 3.1.1 | Test Cases and Data Collecting | 31 |
| 3.2 | Procedure for Testing | 32 |
| 3.2.1 | Testing Procedures | 33 |
| 3.3 | Entra ID Setup | 37 |
| 3.3.1 | Smart Lockout | 37 |
| 3.3.2 | Multi-Factor Authentication | 38 |
| 3.3.3 | Creating Test-Users | 38 |
| 3.4 | Microsoft Sentinel Deployment | 38 |
| 3.4.1 | Microsoft Entra ID Setup | 39 |
| 3.4.2 | Microsoft Entra ID Protection Setup | 40 |
| 3.4.3 | How the Setup and Configuration is Done | 40 |
| 3.4.4 | Testing | 40 |
| 3.5 | Setup of Conditional Access Policies | 42 |
| 3.5.1 | Tools used | 42 |
| 3.5.2 | Setup and Configuration | 43 |
| 3.6 | Evaluation of The Chosen Method | 43 |
| 3.6.1 | Performing The Attack Techniques | 44 |
| 3.6.2 | Entra Setup | 44 |
| 3.6.3 | Sentinel Setup | 45 |
| 3.6.4 | Summary | 45 |
| 4 | Results | 47 |
| 4.1 | Outline of Chapter | 47 |
| 4.2 | Valid Account | 47 |
| 4.2.1 | Valid Account: New location | 47 |
| 4.2.2 | Valid Account: New device | 54 |

- 4.2.3 Valid Account: Atypical Travel 57
- 4.3 MFA Request 59
 - 4.3.1 MFA Request (explicit deny) 59
 - 4.3.2 MFA Request (no answer) 63
- 4.4 Gather Victim Info 65
- 4.5 Proxy 66
 - 4.5.1 Proxy (without MFA) 66
 - 4.5.2 Proxy (with MFA) 69
- 4.6 Brute Force 71
 - 4.6.1 Brute Force: Password Guessing 71
 - 4.6.2 Brute Force: Password Spray 72
- 4.7 Steal Web Session Cookie 75
 - 4.7.1 Steal Web Session Cookie (same IP address) 75
 - 4.7.2 Steal Web Session Cookie (new IP address) 77
- 4.8 Account Manipulation 79
 - 4.8.1 Add Cloud Roles 79
 - 4.8.2 Add Cloud Credentials 81
- 4.9 Create Account 81
- 5 Discussion 83**
 - 5.1 Outline of Chapter 83
 - 5.2 Sub-Research Question 1 83
 - 5.3 Sub-Research Question 2 88
 - 5.3.1 Reducing False Positives 88
 - 5.3.2 Reducing False Negatives 89
 - 5.3.3 Conclusion 92
 - 5.4 Sub-Research Question 3 93
 - 5.5 Sub-Research question 4 95
 - 5.6 Main Research Question 96
- 6 Conclusion 101**
 - 6.1 Answering the Problem Statement 101
 - 6.2 Project Plan and Goals 102
 - 6.2.1 Effect and Project Goals 102
 - 6.2.2 Gantt Chart 102
 - 6.3 Further Research 103
 - 6.3.1 Areas for Improvement 103
 - 6.3.2 Conclusion and Future Directions 103
- Bibliography 105**
- A Additional Material 115**
 - A.1 Script for Creating Users and Groups 117
 - A.2 Microsoft Entra ID - solution 124
 - A.3 Microsoft Entra ID Protection - solution 126
 - A.4 MFA explicitly deny, full table 127
 - A.5 Scripts used to perform the brute force attacks 129
 - A.5.1 Brute Force: Password Guessing 129

| | | |
|----------|-------------------------------------------------------|------------|
| A.5.2 | Brute Force: Password Spraying | 129 |
| A.6 | User CSV File | 129 |
| B | Standard Agreement | 131 |
| C | Project Plan | 139 |
| D | Gantt Chart | 149 |
| E | Time Table | 151 |
| F | Minutes of Meetings with Supervisors | 153 |

Figures

| | | |
|------|------------------------------------------------------------------------------------------------------------------------------------------|----|
| 3.1 | Overview over the incident dashboard in Microsoft Sentinel | 41 |
| 3.2 | Overview over specific incident, with the <i>Events</i> marked | 41 |
| 3.3 | Overview over some of the query and table | 42 |
| 3.4 | Which Conditional Access policies are configured from the Zero-Trust template, see section 2.4.3 | 43 |
| 4.1 | From Microsoft Entra Identity Protection (ID Protection) risk detections: Risk detection registered | 49 |
| 4.2 | From ID Protection risk detections: Details of triggered risk detection. | 50 |
| 4.3 | From Microsoft Entra ID (ME-ID) sign-in logs: User trying to sign in from Egypt, but failing Multi-Factor Authentication (MFA) challenge | 50 |
| 4.4 | From ME-ID sign-in logs: Activity details sign-ins, user is unable to sign in due to failed MFA challenge | 51 |
| 4.5 | From ME-ID sign-in logs: Activity details, Conditional Access Policy (CAP)s applied | 51 |
| 4.6 | From ME-ID Sign-in logs: The four CAPs that were applied | 57 |
| 4.7 | From ME-ID Sign-in logs: The three CAPs that were applied | 58 |
| 4.8 | From ME-ID Sign-in logs: All the CAPs that were applied after denying the MFA challenge | 62 |
| 4.9 | From ME-ID Sign-in logs: All the CAPs that were applied after completing the MFA challenge | 62 |
| 4.10 | From ME-ID Sign-in logs: All the CAPs that were applied after denying the MFA challenge | 64 |
| 4.11 | From ME-ID Sign-in logs: All the CAPs that were applied after completing the MFA challenge | 65 |
| 4.12 | .txt file uploaded to a public GitHub repository containing user credentials (see 3.2.1). | 66 |
| 4.13 | Risky User (without MFA) detected in ID Protection, after risk has been dismissed by admin. | 67 |
| 4.14 | From ME-ID Sign-in logs: All the CAPs that were applied and their result status | 68 |
| 4.15 | From ME-ID Sign-in logs: Authentication method failing | 68 |
| 4.16 | User attempting to sign in using proxy without MFA | 68 |
| 4.17 | Risky User detected in ID Protection (with MFA) | 70 |

| | | |
|------|----------------------------------------------------------------------------------------------|-----|
| 4.18 | From ME-ID Sign-in logs: Authentication method being succeeded . | 70 |
| 4.19 | From ME-ID Sign-in logs: All the CAPs that were applied and their result status | 71 |
| 4.20 | From ME-ID Sign-in logs: All the CAPs that were applied | 75 |
| 4.21 | From ME-ID Sign-in logs: Authentication method being previously satisfied | 76 |
| 4.22 | From ME-ID Sign-in logs: All the CAPs that were applied and their result status | 76 |
| 4.23 | From ME-ID Sign-in logs: Authentication method being previously satisfied | 77 |
| 4.24 | From ME-ID Sign-in logs: All the CAPs that were applied and their result status | 78 |
| 4.25 | From ID Protection Risk Detection Details: Details of the risk detected | 78 |
| 5.1 | Sentinel overview dashboard | 84 |
| 5.2 | ID Protection overview dashboard | 84 |
| 5.3 | Sentinel incident dashboard | 85 |
| 5.4 | ID Protection risk detections dashboard | 85 |
| A.1 | Password Guessing Script | 129 |
| A.2 | Password Spraying Script | 129 |

Glossary

Audit logs Audit logs keep track of changes made to the system and who made them, when and where they happened, and what happened as a result [1]. 10

Autonomous System Number An Autonomous System Number is an IP address prefix used for quicker routing of packets between Internet Service Providers (ISPs) [2]. xiii, 13

Bruteforce A brute force attack is identified by the attacker attempting to sign in to a single user by using an extensive list of passwords. 14, 124, 125

Password spraying A password spraying attempt is identified by an attacker attempting sign-in to a wide array of users using a small sample of passwords, or just one. 13, 14

Provisioning logs Provisioning Logs are diagnostic logs that track actions of a provisioning service [1][3]. 10

Sentinel Sentinel is a Security information and event management (SIEM) and Security orchestration, automation, and response (SOAR) solution provided by Microsoft. xii, 5, 22, 25, 26, 31–33, 38, 40, 41, 43, 45, 47, 50, 51, 54, 57–59, 62, 63, 66, 69, 71, 72, 75, 77, 79–81, 83–99, 101–103, 124

Sign-in logs In ME-ID all sign-ins are logged and the data is then used to monitor risky sign-ins, provide insight into application usage, and more [1]. 10

UEBA *User and Entity Behavior Analytics* is a method used by Microsoft Sentinel to build behavioural profiles with the help of different techniques and machine learning, with the goal of identify and determine compromised entities through abnormal activity [4]. 40, 90, 96, 125

Acronyms

- AI** Artificial Intelligence. 23
- API** Application Programming Interface. 15
- Azure AD** Azure Active Directory. 9
- CA** Conditional Access. 5, 10, 18–20, 22, 31, 38, 43, 44, 47, 50, 51, 58, 62, 65, 69, 71, 72, 75, 76, 78, 81, 82, 91–96, 98, 99, 124
- CAP** Conditional Access Policy. xi, xii, 2, 18–21, 31, 42–45, 50, 51, 57, 58, 61, 62, 64–68, 70, 71, 74–78, 80, 89, 91, 93–95, 97–99, 101, 102, 104
- Content Hub** Microsoft Sentinel Content Hub. 38, 40, 45, 87, 98
- Graph** Microsoft Graph. 15
- IAM** Identity access management. 1, 2, 5, 7–9, 12
- ID Protection** Microsoft Entra Identity Protection. xi, xii, 1, 2, 5, 10, 12, 15–17, 19, 31, 32, 34, 43–45, 47, 49, 50, 56, 58, 61, 62, 64–67, 69–72, 74–78, 80–88, 91–99, 101–103
- IoT** Internet of Things. 8
- ISPs** Internet Service Providers. xiii
- ME-ID** Microsoft Entra ID. xi–xiii, 1, 2, 5, 9–12, 16, 31, 35, 38, 43, 45, 50, 51, 57, 58, 62, 64–66, 68, 70, 71, 75–78, 92, 97–99, 101, 124
- MFA** Multi-Factor Authentication. xi, 8, 10, 16, 19–22, 27, 28, 31, 33–38, 44, 45, 47, 49–51, 54, 56–59, 61–65, 67–71, 74–77, 89–94, 98, 124, 125
- OOTB** Out-of-the-box. 25
- OS** Operating System. 22
- SAML** Security Assertion Markup Language. 13

SEM Security event management. 22

SIEM Security information and event management. xiii, 1, 22, 23, 26, 33

SIM Security information management. 22

SOAR Security orchestration, automation, and response. xiii, 22, 23, 26, 33, 97, 103

SOCs Security Operations Centers. 22

SSO Single Sign-On. 8, 92

VM Virtual Machine. 16, 17, 44

VPN Virtual Private Network. 13, 14, 16, 17, 36, 44, 91

Preface

We want to thank TietoEvry, represented by Thor Larsen, for making this thesis possible and providing us with a team with technical competence. We would also like to give a special thanks to our supervisor from NTNU, Tor Ivar Melling, for listening to our silly rambling, helping us evaluate our own text and giving us valuable guidance and support.

We want to express our gratitude to Tobias Moe, Catherine Kjørsvik and Peder Markus Nielsen for their guidance and expertise throughout the process. They have contributed by directing us on the right path, which led to the thesis we have today.

Chapter 1

Introduction

1.1 Background

The use of cloud resources in businesses today is widespread. In Norway, 71% of private enterprises use some form of cloud computing in their business activities, excluding financial activities. This widespread use of cloud resources has more than doubled since 2014 and has the potential to grow even more in the future [5]. To manage access to these resources, many businesses have had to generate more and more identity-related information.

An international report from 2022 found 67% of businesses to be experiencing *identity sprawl*, a rapid growth in the amount of identity-related information a business needs to work with. Businesses generally found Identity access management (IAM) to be of vital importance, but a majority of 61% said that the management process was too expensive, where 66% of them cited technical debt as the reason. When asked about identity-related breaches, 84% of the respondents said they had experienced one, and 67% had one in 2021 [6]. Many of these breaches could have been prevented if the businesses had performed proper IAM without roadblocks such as technical debt and costs.

To build on the need for secure IAM, the Norwegian Police Security Service (PST) presented in their risk report for 2023 that they suspected a rise in attacks targeting people directly, instead of business infrastructure. The reason for targeting people is to get to their users, which could be used as a pivot for further escalating the stakes of the attack [7]. A report from IBM in 2023 shows that the average cost of a data breach for a Scandinavian business was 1,91 million dollars [8]. The potential damages that could incur from poor IAM could lead to massive economic damages to a business. This shows the importance of securing IAM within the still-growing cloud environments of Norwegian enterprises. There exist many cloud providers, but the most used cloud platform in Norway is Azure[9].

Microsoft's cloud platform *Azure* provides customers with a solution for IAM within the cloud through ME-ID for identity management, and ID Protection for management of access and security. Microsoft also provides a SIEM solution in

Azure with Sentinel, which can be used to manage security incidents as they arise. Sentinel also has automation capabilities that enable businesses to configure automated responses to incidents and have queries that look through multiple incidents and try to find any correlations between them.

In 2021 Microsoft launched *Content Hub*, a marketplace where organisations could sell rulesets they had made for Sentinel [10]. Here, Microsoft also provides rulesets for connecting ME-ID and ID Protection into Sentinel. These rules are available for free and are made to be easy to set up.

1.2 Problem Statement

In this thesis, we will investigate the rulesets provided by Microsoft in Content Hub for the IAM solution ME-ID and the connected security suite ID Protection, to find out how effective the solution is at detecting incoming attacks against users. As most businesses report that technical debt is an issue for proper IAM, we are choosing to approach the security setup by following Microsoft's recommendations and providing ready-to-go solutions to simulate how an organisation without high-skilled employees could construct a security setup for IAM in Azure. This means that we will focus on the enabling of pre-made rulesets in Content Hub for ME-ID and ID Protection and best-practice setup of CAPs within ID Protection.

We will also attempt to discover which rules need to be added, or changed, to lower the amount of false positive and false negative results. However, this should not come at the cost of any true positive result as this could lead to an attack not being detected.

To evaluate the setup of Sentinel and ID Protection, we will utilise a practical approach involving the simulation of attack techniques observed in real data breaches. To assess the comprehensiveness of our testing, we are referencing the framework MITRE ATT&CK and Microsoft's guide for simulating risk detections. Microsoft's guide only covers a few risk detections, which will serve as a baseline for our evaluation. Any additional risk detections triggered during our testing will indicate the thoroughness of our testing.

1.3 Research Questions

The main research question for our thesis is the following:

Main research question *How well do the rulesets provided by Microsoft in Sentinel Content Hub for Entra ID and Entra ID Protection with a best-practice setup of Conditional Access policies secure an organisation against user identity-based threats?*

This question could be further broken down into four specific sub-questions:

1. *How does Sentinel, configured with rulesets for Entra ID and Entra ID Protection provided by Microsoft in Content Hub, provide any additional security features which are not available through the Entra ID Protection dashboard?*
2. *How can the rulesets in Sentinel be modified to reduce the number of false positive and false negative results?*
3. *How does the use of best-practice Conditional Access policies affect what is detected while using Microsoft's rule-sets from Content Hub and risk detection in ID Protection?*
4. *Can we trigger any additional risk detections for user identities beyond those presented by Microsoft?*

1.4 Effect Goals

With these research sub-questions, we hope to achieve the following effect goals:

1. Assess the security provided to user identities by a setup which only follows Microsoft's recommendations.
2. Show how a practical approach to testing a business' security setup with Entra ID Protection and Sentinel can be performed and assessed.

1.5 Project Goals

When finishing the project, we want to have achieved the following:

1. Have tested the rules provided in Content Hub by Microsoft for Entra ID and Entra ID Protection.
2. Discovered what limits and difficulties are present when attempting to simulate attacks.
3. Collected a guide for best-practice setup for Conditional Access policies in Entra ID Protection according to current guidelines.
4. Have tested what differences the use of best-practice Conditional Access policies has on the incident detection and handling in the Sentinel setup.
5. Found what more a business will be able to see and do when using Sentinel compared to using the Entra ID Protection dashboard.

1.6 Partner Organisation

This thesis was completed in cooperation with the IT consultant company TietoEvy. They suggested the topic of which we were to study, but we chose the specific problem statement and the focus on ready-to-go solutions. The choice of Azure as the cloud platform to use in this thesis was due to TietoEvy's expertise in Microsoft Azure.

While working on the thesis they provided guidance and help. This included both technical help and proofreading.

1.7 Thesis Outline

Our thesis is structured into three parts, with the largest section being the main part, containing six chapters. Each chapter serves a specific purpose in presenting our research findings and analysis. The start of the thesis consists of information directly relevant to the reading of the thesis. The structure is as follows:

- **Abstract:** A brief overview of the purpose of the thesis, the methods used, the main findings, and the conclusions.
- **Sammendrag:** Same content as in Abstract, but written in Norwegian.
- **Contents:** An overview of the thesis's structure and content to help the reader in navigating through the document.
- **Figures:** A list of all figures, tables, and diagrams included in the thesis, along with corresponding numbering and page references.
- **Glossary:** A list of key terms with definitions to assist the reader in understanding specialised terminology used in the thesis.
- **Abbreviations:** A list of abbreviations and acronyms used in the thesis, along with their full meanings.
- **Preface:** A brief introduction describing the background of the thesis, the motivation for carrying it out, and gratitude to any contributors or supporters.

In the middle, we have the main part of our thesis. This part encompasses all details on the tests conducted and their results. This is the structure:

- **Chapter 1 - Introduction:** In Chapter 1, we introduce the purpose of our thesis and articulate the objectives we aim to achieve. We will also provide some background for the thesis.
- **Chapter 2 - Theory:** Chapter 2 delves into the theoretical framework; exploring necessary concepts to provide a foundational understanding of how they work.
- **Chapter 3 - Method:** The third chapter outlines the methodology employed to address our research questions, detailing the data collection, analysis techniques, and evaluation methods utilized. We will also reflect on the strengths and limitations of our methodology.
- **Chapter 4 - Results:** During chapter 4 we present the results obtained from our testing procedures, including any identified findings, trends, or patterns.
- **Chapter 5 - Discussion:** Chapter 5 engages in a comprehensive discussion, contextualizing our results within the broader scope of the thesis's objectives, theoretical underpinnings, and previous research.
- **Chapter 6 - Conclusion:** Finally, in Chapter 6, we summarize our findings and draw conclusions based on our research questions, offering insights and potential directions for future work and investigation.

Then at the end of the thesis:

- **Bibliography (Bibliografi):** A list of all sources and references used in the

report, formatted in the IEEE referencing style.

- **Appendix:** Supplementary material containing extra information, not essential to the main text but still relevant to the thesis.

1.8 Scope and Delimitation

The scope of this thesis is defined by the conditions outlined in our problems statement (1.2). Our primary goal is to assess whether the security measures, implemented per Microsoft's recommendations for ME-ID, ID Protection and Sentinel, can detect potential attacks relating to user identities. Therefore, our focus will predominantly be detecting attacks targeting user sign-in activity within a cloud environment.

However, it is important to note certain limitations within our study. Firstly, the project will not involve the simulation of attacks originating from known malicious IP addresses or the deployment of actual malware. Additionally, our thesis will not use or extend to the configuration and utilisation of tools or platforms other than those detailed in our problem statement. Consequently, some of the risk detections in ID Protection and some rules in Sentinel fall outside the scope of our thesis.

The set-up process of Conditional Access (CA) is included as part of this thesis. As described in the problem statement, these policies will be set up following Microsoft's best-practice recommendations. These policies are often customised to suit each organisation. Similarly, a customised setup of user identity architecture within Azure is crucial for proper IAM. However, this level of tailoring is beyond the scope of our thesis.

A cloud environment is never cost-free. We operated under an E5 Security subscription to Azure. With this subscription, we can include all relevant functions within Sentinel and ID Protection in our thesis. However, assessing whether an organisation should opt for this specific subscription or another falls outside our scope.

By scoping our thesis in this manner, we enhance our understanding of each specific response elicited by the setup. A security setup with fewer distinct components will facilitate a more precise evaluation of how Microsoft's recommended setup works as an out-of-the-box solution. Furthermore, adopting a minimalist configuration of Azure around this security setup will also render our conclusion more general, enabling a broader audience to utilise our results as a foundation for configuring a security setup tailored to their organisation's needs.

Chapter 2

Theory

2.1 Identity Access Management

2.1.1 What is Identity Access Management

Identity access management (IAM) ensures secure access to an organisation's resources, including emails, databases, data, and applications, for verified entities with minimal interference. It is crucial in today's hybrid work environment, to facilitate controlled access and ensure that sensitive data and functions are restricted to authorised individuals only. This system is foundational to an organisation's cyber security, managing and verifying access attempts efficiently and securely [11] [12].

Identity

According to Microsoft, a digital identity is a set of distinct identifiers or attributes that represent a human, software component, machine, asset or resource in a computer system, such as [13]:

- An email address
- Sign-in credentials (username/password)
- Bank account number
- Government-issued ID
- MAC address or IP address

These identities serve the purpose of authenticating and authorizing access to different resources, communicating with other individuals, conducting transactions, and other purposes [13].

There are three types of identities [13]:

User identities are people such as internal employees (both administrative staff and frontline workers) as well as external users (including customers, consultants, vendors, and partners).

Workload identities refers to software entities such as applications, services, scripts, or containers.

Device identities represent physical devices like desktop computers, mobile phones, Internet of Things (IoT) sensors, and IoT managed devices. It's important to note that device identities are distinct from human identities.

2.1.2 Authentication and Authorisation

Authentication and *authorisation* are fundamental components of IAM. They serve different purposes in the access control process [13]:

Authentication is the process of verifying the identity of users, software, components, or hardware devices through credentials such as usernames and passwords, biometrics, or security tokens [13]. MFA adds a layer of security by requiring multiple forms of evidence for identity verification, while Single Sign-On (SSO) simplifies the authentication experience by allowing users to authenticate once. IAM then acts as the source of truth for user identities across multiple resources, minimising the burden of signing in to multiple systems separately [13].

Authorisation determines access to resources based on verified identities, following authentication. While authentication validates identity, authorisation regulates resource access according to predefined permissions and policies [13].

Authentication and authorisation are often misconstrued as interchangeable terms by users because they experience them as a single process. However, they are two separate processes: *authentication* confirms identity, while *authorisation* controls resource access [13].

2.1.3 How Identity Access Management Works

IAM operates on two main principles: *identity management* and *access management*.

Identity management involves the authentication of a user's identity against an identity management database, which must be constantly updated. This process may include MFA for added security [11].

Access management is the second component and manages what resources a verified user can access, ensuring that users only have access to necessary resources based on their roles [11].

This dual approach ensures that authentication and authorisation occur securely and accurately with every access attempt [11].

2.1.4 Why Identity Access Management Is Important

IAM is a critical component of an organisation's cyber security strategy for several reasons. It helps balance the accessibility of sensitive data and resources, ensuring they are accessible to authorised individuals while preventing unauthorised access. With cyber threats constantly evolving, IAM plays a crucial role in defending against attacks and minimising the impact of data breaches [11]. Furthermore, IAM supports compliance with various regulatory requirements by automating the auditing process and managing data access governance efficiently. Implementing IAM not only strengthens an organisation's security posture but also enhances operational efficiency and regulatory compliance [11].

Identity Access Management and Cloud Computing

IAM is a crucial component in cloud computing, addressing the limitations of traditional security measures such as usernames and passwords. Normal usernames and passwords are no longer strong enough to protect against breaches. Passwords, being vulnerable to hacking, sharing, or being forgotten, cannot provide the robust security that is required for today's organisational needs [11].

IAM systems address the challenges of monitoring and managing access attempts manually, by making the management of identity attributes possible, enabling organisations to grant or restrict access based on roles efficiently, and providing mechanisms to flag anomalies and security breaches promptly. This capability is crucial in cloud computing, where it demands agile and robust security measures to protect sensitive data and resources from unauthorised access [11].

2.2 Microsoft Entra ID

2.2.1 What Is Microsoft Entra ID?

Microsoft Entra ID (ME-ID) is what was formerly known as Azure Active Directory (Azure AD). ME-ID is an integrated cloud identity and access solution, and it provides a single place to store information about digital identities. More specifically, ME-ID is an IAM solution (see 2.1) from Microsoft that helps organisations secure and manage identities for hybrid and multi-cloud environments [14] [15].

Microsoft Online business services like Microsoft 365 or Microsoft Azure use ME-ID for sign-in activities and to help protect other identities. By subscribing to any of these services you automatically get access to ME-ID Free. However, to get access to more features and enhance the Microsoft Entra implementation you'll need to get a paid licence. This licence will provide self-service, enhanced monitoring, security reporting, and secure access for mobile users [16].

2.2.2 Features in Microsoft Entra ID

Based on the specific licence you choose you get access to some or all of the following features [16]. Due to the scope outlined in section 1.8, only the relevant features will be described:

Authentication Verifying, or *authenticating*, credentials when a user signs in to a device, application, or service, is one of the main features of an identity platform. In ME-ID authentication is more than just verification of a username and password. ME-ID authentication includes the following component:

- Self-service password reset
- Microsoft Entra MFA
- Hybrid integration to write password changes back to on-premises environment
- Passwordless authentication

Conditional Access CA is a security feature by Microsoft that lets organisations decide who can access what based on certain conditions. These conditions can be the user's identity, device, location, or the security of the network they are connected to [17]. For more information, see 2.4.

Enterprise users With ME-ID provides management services to organisations, allowing them to assign licences, manage groups and users, and add or manage domain names [18].

Hybrid identity Microsoft's identity solutions cover both on-premises and cloud-based capabilities. These solutions establish a shared user identity for accessing resources, no matter where they are located. This concept is known as hybrid identity [19].

Identity Protection ID Protection gives organisations the ability to detect, investigate, and remediate identity-based risks. Furthermore, suspicious actions can be responded to with configured policies, and then organisations can take appropriate action to resolve them [20] [16]. For more information, see 2.3.2.

Monitoring and health Microsoft Entra monitoring and health enables organisations to gain insight into the security and usage patterns in their environment. Activity logs in ME-ID are divided into three types; Audit logs, Sign-in logs, and Provisioning logs [16] [21].

2.2.3 Concepts in Microsoft Entra ID

Users

Within a Microsoft Entra tenant, you can either be an internal or external user. This makes it so there are 4 types of users [22]:

Internal member Users that are full-time employees of the organisation

Internal guest Users that have an account in the organisation's tenant, but only have guest-level privileges.

External member Users that have member access to the organisation's tenant, but authenticate by using an external account.

External guest Users who have guest-level privileges and who authenticate using an external method.

Groups

When creating new groups in ME-ID there are two group types you can choose from:

Security Security groups contain users, devices, groups, and service principals as their members. It is the users and service principals who are the owners of this group. With security groups, you can apply licenses to users based on their group membership. You can also apply a security policy to these groups to grant a set of permissions to all the members at once, instead of having to do this to each member individually. A security group also enables you to give people outside of the organisation access to the group [23].

Microsoft 365 Microsoft 365 groups can only have users as their members. These groups are used to ensure that groups of people have consistent permissions to a group of related resources and allow you to set up a collection of resources to share [23].

Then there are different types of memberships to choose for the members [24] :

Assigned Users are manually added to become members of a group.

Dynamic user Users are automatically added or removed as members of a group by using dynamic membership rules.

Dynamic device Devices are automatically added or removed as members of a group by using dynamic group rules.

2.2.4 Azure Smart Lockout

Azure Smart Lockout, or Smart Lockout, is a setting in ME-ID which blocks a user from signing in after several failed attempts and is always turned on. The default number of failed sign-ins needed for Smart Lockout to lockout a user in an Azure Public Tenant is 10 attempts. After the initial lockout, the user will continue to be locked out for each subsequent failed sign-in attempt. Only failed sign-in attempts which are different from the previous three previous attempts will lead to a lockout, and familiar locations will have different lockout-timers to an unfamiliar location [25].

2.3 Microsoft Entra Identity Protection

This section will present how Microsoft Entra ID Protection works and what it consists of. We start with a brief explanation of how ID Protection functions on a superficial level before delving into how specific risk detections work. Afterwards, we proceed to how the suspicious events are presented to a security employee tasked with supervising identity-related threats before presenting methods of simulating a real attack by manually triggering risk detections. The conditional access part of ID Protection will be discussed in section 2.4.

2.3.1 What Is It and What Does It Do?

ID Protection is a security product for securing the Azure cloud native IAM solution ME-ID. The security suite delivers surveillance of the IAM over the Azure tenant making it possible to detect, investigate and remediate potentially malicious activity, powered by machine learning [20].

2.3.2 Detecting Risk

The events which ID Protection detects as suspicious are called *risk detections* [20]. These detections are designed to detect specific MITRE ATT&CK pattern types [26]. However, not all potentially suspicious activity leads to a risk detection.

A risk detection is only generated by ID Protection when there is a real chance for a user to be compromised. One of the requirements is that the correct credentials are in use for the user [27]. A risk detection does not always end up increasing the perceived risk for a given user. The detection can also be confirmed as a false positive by a security admin or by ID Protection itself at a later stage. The risky event identified by the detection could be remediated through access policies at a later stage (see 2.4.1) [27].

Two Types of Risk

There are two types of risk detections, *real-time* and *offline detections*, which show at what time the different risk detections are analysed [27]:

Real-time detection A risk detection marked as *real-time* is triggered procedurally as new events are registered. Risk detections of this type are usually registered within 5 to 10 minutes.

Offline detection An offline risk detection is analysed at a later time after the suspicious event occurred. A detection of this type might take around 48 hours to be registered by ID Protection.

Risk Detection for User Identities

The risk detections which trigger for user-identities are divided into two categories. The first category listed here is risk detections which trigger when detecting specific sign-in events, and the second category, user risk detections, trigger on a given user. The detection type of each risk detection is noted in parentheses after the name. Due to the scope outlined in Chapter 1 (1.8), only risk detections within the defined scope will be described.

Sign-in risk detections A sign-in risk detection triggers when detecting a specific sign-in event. This signals that a given event is suspicious and that it might be an adversary which performed the attempted sign-in. Important to note again is that the first credentials need to be correct for a risk detection to occur [27].

Atypical travel (offline) This risk detection indicates that there might be an adversary trying to sign in using the same credentials from a distant location at the same time as the owner of the user. For this detection to trigger, two sign-in attempts must have been registered from different geographically distant locations during a short period. The algorithm behind this detection ignores false positives such as the use of a Virtual Private Network (VPN) and IP addresses often in use by users in the organisation. For this detection to trigger, the algorithm needs a learning period of 14 days or 10 logins [27].

Anomalous token (real-time and offline) This detection triggers when unusual characteristics are discovered in the Security Assertion Markup Language (SAML) token. It implies that an attacker might have attempted signing in by reusing an old session or refresh token, or that a token is coming from an unfamiliar location which the token was not issued for. Microsoft notes that the anomalous token risk detection is more likely to create false positives by design to increase the likelihood of detecting replayed tokens [27].

Unfamiliar sign-in properties (real-time) This detection triggers when sign-in activity is registered for a user with sign-in properties, such as IP address, browser, device, tenant IP subnet, location or Autonomous System Number, which has not been seen before. Due to this detection relying on past sign-in history, there has to be a learning period of a minimum of five days. If the user has long periods of inactivity the user will go back into learning mode [27].

Password spray (offline) This detection detects when a Password spraying attack is registered. This is when a wide sample of users are attacked using a small sample of passwords in the same method in hopes of having a successful sign-in [27].

Anonymous IP address (real-time) This risk detection triggers when there is a sign-in attempt which is registered from an anonymous IP address.

Services which could trigger this alarm are the use of the Tor browser¹ or an anonymous VPN service [27].

User risk detections The other category of risk detections is those related to users. These are risk detections which not necessarily link to a specific sign-in event, but still pose a risk for the user to be compromised [27].

Leaked credentials (offline) When this risk detection triggers, it indicates that a user's credentials have been leaked online. Microsoft regularly performs an investigation online which looks for potentially leaked credentials [27]. This risk detection is also present for workload identities [29].

Risk Detections and MITRE Attack Patterns

As described at the beginning of section 2.3.2 each risk detection is designed to detect specific MITRE ATT&CK patterns. The following section presents the relevant attack patterns alongside the specific risk detections designed to identify each pattern² [26]:

Access using a valid account: Cloud accounts (T1078.004) An attack pattern identified by MITRE is the use of existing cloud accounts to gain initial access. After gaining access the use of an existing account might help disguise malicious operations performed by the attacker [30].

- Unfamiliar Sign-in Properties

Account Manipulation (T1098) After an attacker has gained initial access to a system they might need to gain a higher level of privileges or change account settings to continue their attack [31].

- Anomalous User activity [26]

Brute Force: Password Spraying (T1110.003) One of the Bruteforce techniques listed by MITRE is Password spraying. This technique employs a small sample of common passwords which are attempted towards a large sample of users. The goal of this technique is to avoid account lockouts while gaining initial access. *Azure smart lockout (2.2.4)* is designed to stop standard Bruteforce attacks with multiple password guesses towards a single user [32]. The following risk detection is created to detect this technique:

- Password spray

¹The Tor browser is a web browser which anonymises the user through the use of multiple layers of encryption and multiple relays between the browser and destination [28].

²Microsoft does not include all the relevant risk detections in their overview of the mapping between risk detections and attack patterns [26].

Gather Victim Identity Info (T1589.001) The gathering of victim information is a technique often used by attackers during the targeting stage of an attack. These can be obtained in many ways; from scouring the internet to black markets. The result of this technique might be new ways for the attacker to perform reconnaissance or initial access to the target system [33].

- Leaked Credentials

Obfuscation/Access using proxy (T1090) To prevent the system owner from identifying the attacker, an adversary might use a service to avoid a direct line of communication between the attacker and the target. This makes it harder for the victim to find out who was behind the attack [34].

- Anonymous IP address

Steal Web Session Cookie/Token Theft (T1539) Another method of gaining access to internet services is stealing another user's session cookie or token. As cookies are stored on the machine and sent together with other network traffic when going to a site an attacker can gain a copy of the cookie through other attack techniques. After gaining the copy of a cookie the attacker can use it to gain initial access to the related service [35].

- Anomalous Token

2.3.3 Presenting Risk

In ID Protection each user is assigned a risk score indicating the likelihood of compromise. This score is calculated based on the number and severity of suspicious events, and can be *low*, *medium* or *high*. This score can be seen for users in the ID Protection dashboard [27].

The ID Protection dashboard can be used by administrators to see the risk score for a given user, see details about risk detections and the history for registered risky events, which can be those who are active at the moment, have had remediations steps or where the risk event has been dismissed. It is also possible for admins to invoke remediating actions such as marking users as confirmed compromised, resetting the password for a user, dismissing a risky event, blocking the user from signing in or gaining further insight into risk detections through associated sign-in events for a user. These steps can also be performed through the Microsoft Graph (Graph) Application Programming Interface (API) [26].

2.3.4 Simulating Risk

Microsoft provides four methods which could be used to simulate an attack. Each of these methods goal is to trigger a specific risk detection [36]. Microsoft claims that these are the risk detections which can be manually triggered, due to the machine learning in use to weed out false alarms. This creates a discrepancy in the

number of risk detections which are relevant for this thesis and the ones Microsoft presents simulation methods for.

Anonymous IP address To simulate this risk detection one would need an anonymisation service. Microsoft suggests using the Tor browser and a test account which has not been registered with MFA. The risk detection will show up in the ID Protection dashboard after 10 to 15 minutes after performing the following [36]:

1. Open the Tor browser.
2. Navigate to <https://myapps.microsoft.com> and sign in using the sign-in credentials of the test user.

Unfamiliar sign-in properties There are many sign-in properties related to a given sign-in attempt. This risk detection can therefore be triggered in two different ways: by simulating a new location or a new device. For either of these, there is a need for a user with at least a 30-day sign-in history, which is registered for MFA in ME-ID. To simulate a new location Microsoft suggests the following:

1. Create a new VPN connection.
2. Navigate to <https://myapps.microsoft.com> and sign in using the sign-in credentials of the test user and fail the MFA challenge.

To simulate a new device Microsoft suggests the following :

1. Create a new Virtual Machine (VM)³.
2. Start the VM and open a browser window.
3. Navigate to <https://myapps.microsoft.com> and sign in using the sign-in credentials of the test user and fail the MFA challenge.

Each of these risk detections will create events in the ID Protection dashboard within 10 to 15 minutes [36].

Atypical travel This risk detection is considered difficult to manually trigger by Microsoft due to the algorithm in use, and they recommend attempting to trigger this risk detection on multiple users for the detection to occur. Microsoft states that you need a user with a sign-in history of 14 days or 10 logins, a method of changing your IP address and a method of changing your user agent⁴. To trigger the detection, perform the following:

1. Open up a browser.

³A virtual machine is a machine which runs as software on another machine. A virtual machine runs its operating system which is separate from the host machine [37].

⁴The user agent used in a web browser is a string of text meant to identify a program which represents a person. In this context it is connected to the web browser [38].

2. Navigate to `https://myapps.microsoft.com` and sign in using the sign-in credentials of the test user.
3. Change your IP address. Microsoft suggests multiple methods such as using a VPN, using a Tor add-on or using a VM set up in a distant location.
4. After changing your IP address, navigate to `https://myapps.microsoft.com` and sign in again, using the sign-in credentials of the test user.

The event will be detected and can be found in ID Protection dashboard within a few hours [36].

Leaked credentials (for workload identities) This is the only risk detection that Microsoft suggests can be simulated for workload identities. To simulate this detection Microsoft suggests the use of Github⁵ for uploading the credentials. The requirements for simulating the risk detection are therefore a Github account, an admin user with at least Security Administrator privilege, and performing these steps [39]:

1. Using the admin user, sign in to the Microsoft Entra admin centre and go to the page *Identity > Applications > App registrations*.
2. Register a new application by selecting *New registration*.
3. Go to *Certificates & Secrets > New client Secret* and create a new client secret for the newly registered application.
4. Write down the value of the secret. This is important as the secret can not be retrieved again later.
5. Find the Tenant ID and Application client ID in the *Overview* page and record them.
6. Disable the application by setting *Enabled for users to sign-in* to **No**, in the page *Identity > Applications > Enterprise Application > Properties*.
7. Using your Github user, create a public repository.
8. Create a file with the ".txt" extension and add the following [36]:

```
"AadClientId" : " < insert - client - id > ",
"AadSecret" : " < insert - application - secret > ",
"AadTenantDomain" : " < insert - domain > .onmicrosoft.com",
"AadTenantId" : " < insert - tenant - id > "
```

9. Commit the file and make sure to push the change to the repository stored at Github.

The risky event will be detected within 8 hours [36].

⁵Github is an online platform where developers can save their work, collaborate and share their code if wanted [39].

2.4 Conditional Access

Conditional Access operates similarly to *if-then* statements in various programming languages, where certain conditions must be met for access to be applied. In its simplest form, it can be explained as; *If* an assignment is met, *then* apply the access conditions. According to Microsoft, these policies can be designed to grant access, limit access with session controls (i.e., session controls can be used within a CAP to enable limited experiences within specific cloud applications), or block access [17].

CAPs are used to apply the right access controls to secure your organisation. To determine the access, CA uses identity-driven signals (e.g. users, groups, IP addresses, devices, applications, real-time risks). It combines these signals to automate decisions and enforce organisational access policies for resources. These CAPs help balance security and productivity, enforcing security controls when needed and staying out of the user's way when not [17][40].

Effective implementation of CAPs requires a detailed understanding of an organisation's security posture, access patterns, and the specific needs of different user groups or personas. According to Microsoft, policies should be tested in report-only mode before full deployment to gauge their impact and effectiveness [41].

Report-only mode: Report-only mode makes it possible to evaluate the impact of CAPs before they are enabled. This means that during sign-in the policies are not enforced. It is also important to note that this does not apply to items included in the *User Actions* scope. Based on Microsoft's documentation, there are four possible result values when a policy in report-only mode is evaluated [42]:

Report-only: Success All configured policy conditions, required non-interactive grant controls, and session controls were satisfied [42].

Report-only: Failure All configured policy conditions were satisfied but not all the required non-interactive grant controls or session controls were satisfied [42].

Report-only: User action required All configured policy conditions were satisfied but user action would be required to satisfy the required grant controls or session controls [42].

Report-only: Not applied Not all configured policy conditions were satisfied [42].

However, it is important to note that these policies are only enforced *after* the first-factor authentication (e.g. a password) is completed [17].

2.4.1 Risk-Based Access Control Policies

According to Microsoft, risk-based policies are access control policies that can be applied to protect organisations when a sign-in or user is detected to be at

risk, by being assigned a high risk score. In Microsoft Entra CA there are two risk conditions: *Sign-in risk* and *User risk*. By configuring these two risk conditions and choosing an access control method, organisations can create risk-based CAPs. ID Protection sends the detected risk levels of each sign-in to CA, and if the policy conditions are met then the risk-based policies apply [43].

Sign-in risk-based CAP: Whenever anyone signs-in, hundreds of signals get analysed by ID Protection in real-time. A sign-in risk level is then calculated, representing the probability that the given authentication request is not authorised before it is sent to CA. Administrators can configure sign-in risk-based CAPs to enforce access controls based on sign-in risk, including requirements such as [43]:

- Block access
- Allow access
- Require MFA

If any risks are detected, users get the option to perform the required access control such as MFA to self-remediate and close the risky sign-in.

User risk-based CAP: Just as with sign-ins, ID Protection also analyses signals about user accounts. A risk score based on the probability that the user is compromised is then calculated. ID Protection uses signals such as risky sign-in behaviour or credentials leaks to calculate the user risk level. Administrators can then configure risk-based CAP to enforce access controls on user risk, including requirements such as [43]:

- Block access
- Allow access but require a secure password change

If the user does a secure password change the user self-remediates the user risk and closes the risky user event.

2.4.2 Zero Trust

The Zero Trust security model has an approach to cyber security where it operates on the assumption that no entity, either within or external to the network perimeter, should be completely trusted. This model is based on the principle of *never trust, always verify*, which involves meticulous verification of all access requests, regardless of their origin [44].

Conditional Access and Zero Trust

Within the Zero Trust framework, CA serves an important role in enforcing the *verify explicitly* principle. This is achieved by evaluating access requests against a set of predefined conditions, effectively implementing a dynamic access control mechanism. Utilising identity-driven signals CA enables the automation of

decision-making processes related to organisational access policies. This integration is important in securing organisations' resources while maintaining a balance between security measures and user productivity. When aligned with the Zero Trust model, CA significantly enhances an organisation's security posture by ensuring that access controls are strictly enforced based on real-time assessments [45][46].

2.4.3 Best-Practice Setup

In formulating the optimal setup for CAPs, the reliance on Microsoft's recommendation is both deliberate and strategic. Microsoft, drawing on its extensive experience and insights, has developed a suite of Conditional Access templates. These templates serve as a streamlined approach to the effective implementation of CAPs in line with Microsoft's guidelines, offering robust protection in combination with the most commonly adopted policies across diverse customer demographics and geographical areas [47]. In our thesis, we define *best-practice* as Microsoft's recommendations.

Following these templates allow organisations to move beyond the deployment of arbitrary security measures, embracing instead a meticulously crafted framework of policies designed to fortify defences across a broad spectrum of cyber threats. This approach not only ensures consistency with industry-leading practices, but also takes advantage of the full capabilities of the Zero Trust architecture. Consequently, it enhances the security and operational efficiency of organisational resources, aligning with the highest standards of cyber security [47].

The template, based on Zero Trust, has a list of 14 policies that collectively help support a Zero Trust architecture [47]:

Require MFA for admins: By requiring MFA on accounts that are targeted by attackers for their assigned administrative rights, the risk of those accounts being compromised is reduced [48].

Securing security info registration: This ensures that the process of registration for MFA and password resets is as secure as accessing critical applications. It prevents unauthorised users from registering security information, enhancing overall security by verifying user identity and compliance [49].

Block legacy authentication: Microsoft recommends that organisations block authentication requests using legacy authentication protocols due to the increased risk associated with using these protocols, and instead require modern authentication. According to Microsoft, legacy authentication is a client or network protocol which is incapable or not configured to do modern authentication (e.g. it sends both username and password) [50].

Require MFA for all users: Based on Microsoft's studies, a user's account is 99.9% less likely to be compromised if MFA is used. This highlights the importance

of MFA in security, indicating that the strength of a user's password becomes irrelevant in comparison.

Require MFA for guest access: Requiring MFA for guest access is an important measure to protect resources from unauthorised access. Microsoft recommends using MFA to ensure that external users require more than just credentials for access. This ensures robust protection and access management, maintaining the security integrity of resources [51].

Require MFA for Azure management: According to Microsoft, many organisations that use Azure services manage them from Azure Resource Manager-based tools (e.g. Azure portal, Azure PowerShell, Azure CLI). These tools can provide highly privileged access to resources. To protect these privileged resources, Microsoft recommends requiring MFA for any user accessing these resources [52].

Require MFA for risky sign-ins: A sign-in risk, according to Microsoft, represents the probability that a given authentication request isn't authorised by the identity owner. Only organisations with Microsoft Entra ID P2 licences can create CAPs incorporating ID Protection sign-in risk detections. The users are protected from registering MFA in risk sessions when the sign-in risk-based policy is enabled. If users aren't registered for MFA, then their risky sign-ins are blocked, and they see an AADSTS53004 error [53].

Require password change for high-risk users: Requiring a password change for high-risk users secures potentially exposed accounts and reduces the window for unauthorised access [27].

Block access for unknown or unsupported device platform: This policy blocks users from accessing company resources when the device type is unknown or unsupported. This ensures that only devices with verified security features and updates can access the system, which minimises the risk of breaches from devices that are not compliant [27].

No persistent browser session: This policy protects user access on unmanaged devices by preventing browser sessions from remaining signed in after the browser is closed, according to Microsoft. This reduces the risk of unauthorised access by automatically logging users out after their session ends [54] [55]

Require approved client apps or app protection policies: By requiring approved client apps or app protection policies, organisations can make sure staff can be productive, but also prevent data loss from applications on devices they don't fully manage [56].

Require compliant or Microsoft Entra hybrid joined device or MFA for all users: organisations that use Microsoft Intune can ensure device compliance with

requirements like PIN unlock, encryption, specific Operating System (OS) versions, and preventing jailbroken or rooted devices. This ensures that only secure, authenticated devices and users can access organisational resources, minimising the risk of unauthorised access and data breaches [57] [58].

Require MFA for admins accessing Microsoft admin portals: Microsoft recommends securing access to any Microsoft admin portals (e.g., Microsoft Entra, Microsoft 365, Exchange, Azure). This will help protect vital access points from unauthorised use and credential thefts, ensure that only authorised admins can access sensitive areas, and facilitate secure account recovery [59] [60].

Block access for users with Insider Risk (Preview): Users with abnormal behaviour can be risky to allow just to sign in. Blocking access for users with Insider Risk will prevent access to sensitive information or systems by individuals deemed to pose an insider risk. An insider risk signal is provided to CA to refine access control decisions [61]. Insider risk signals are based on contextual factors like user behaviour, historical patterns, and anomaly detections [62]. Microsoft Purview⁶ needs to be enabled before the signals can be used in CA [61].

2.5 Microsoft Sentinel

Microsoft Sentinel is a cloud-based solution designed to modernise the traditional Security Operations Centers (SOCs). By using Sentinel, users will be provided with access to both SIEM and SOAR functionalities [65]. These functionalities are further used to secure an organisation environment through different key features, which will be explained later on in section 2.5.4.

2.5.1 What Is SIEM?

SIEM is a security solution that combines Security information management (SIM) and Security event management (SEM) into one security management system that helps organisations detect, analyse and respond to potential security threats before they harm or disrupt any business operations [66].

There exist different SIEM solutions, such as Microsoft Sentinel, and their capabilities will vary. However, they generally share the same set of core functions:

Log management SIEM systems collect event/log data from different sources in an organisation's infrastructure, organise it, and then decide if it shows any signs of security threats [67].

⁶Microsoft Purview is a set of tools designed to provide solutions that can help organisations govern, protect, and manage data [63]. Microsoft Insider Risk Management makes it possible to create policies to manage security and compliance, and correlates various signals to identify potential malicious or inadvertent insider risks [64].

Event correlation The data is then sorted, and advanced analytics are used to identify any relationships or patterns, enabling quick detection and response to security threats [67].

Incident monitoring and response SIEM technology monitors security incidents, and often visualises them in a dashboard, where security teams monitor the activity, triage alerts, identify threats, and respond to detected threats [67].

2.5.2 What Is SOAR?

SOAR is a set of services and tools that help organisations automate prevention and response to security threats. Typically, a SOAR solution consists of three core features:

Security orchestration This feature involves how the platform connects and coordinates the hardware and software tools in an organisation's security system [68].

Security automation SOAR can automate tasks, often those that are low-level, time-consuming, and repetitive, allowing them to be executed automatically. This is often accomplished through playbooks or workflows that automatically run when they are triggered by rules or incidents [69].

Incident response The orchestration and automation capabilities of SOAR solutions allow it to serve as a console for security incident response, where metrics and alerts can be aggregated and displayed in a central dashboard [68].

2.5.3 SIEM and SOAR

A natural setup involves using SIEM and SOAR together, as they complement each other. This approach is recommended because SOAR solutions are primarily used to orchestrate and automate threat responses, while SIEM offers greater visibility into activities through threat detection, log management, and more. This means that, in tandem, SIEM will collect and analyze data, while SOAR operates based on the data provided by the SIEM solution. By combining these two components, organisations establish a complete solution for risk detection, visibility, and response, as demonstrated by Microsoft Sentinel [69].

In general, this means that Microsoft Sentinel is a solution that helps organisations uncover and respond to threats effectively. Microsoft Sentinel serves as a comprehensive solution for threat recognition, examination, and response by utilising SIEM capabilities. Overall, Microsoft Sentinel collects log data from configured environments, detects threats and minimises false positives, investigates threats by leveraging Artificial Intelligence (AI), and orchestrates responses with automation through SOAR [65].

2.5.4 How Does Microsoft Sentinel Work?

Microsoft Sentinel collects data through data connectors and then visualises and monitors this log data using workbooks. In Microsoft Sentinel, you can set up different rules and procedures, such as analytics rules, automation rules, and playbooks. These rules help Microsoft Sentinel discover security threats and abnormal behaviour across the environment. Microsoft Sentinel also allows you to understand the cause of your security threats through deep investigation tools and enables you to proactively hunt for anomalies through built-in queries. Microsoft Sentinel also offers the opportunity to use notebooks to extend the scope of what you can do in the solution [70]. The following section will provide a more precise description of each key feature mentioned:

Data connectors Data connectors enable connections between Microsoft Sentinel and other services that you want to collect log data. These connectors ingest the data into Microsoft Sentinel, allowing organisations to gather insight across their environment [71].

Workbooks Once the data sources are connected to Microsoft Sentinel, workbooks are used to visualise and monitor the data. Here data can be represented in different styles, such as charts, graphs and tables, and also filtration of data is possible for a more effective investigation [72].

Analytics rules Analytics rules are customised rules that can help discover threats and abnormal behaviour in an environment. They search for specific events in an environment and alert when certain thresholds or conditions are met [73].

Automation rules Automation rules allow organisations to manage automation in Microsoft Sentinel by allowing them to define and coordinate a set of rules that can be used across different scenarios, such as assigning the scenario to the right personnel or suppressing noisy incidents [74]

Playbooks Playbooks are collections of procedures that can be run in response to incidents⁷, alerts or specific entities. They can help organisations automate and orchestrate responses and can be set to run automatically when attached to an automation rule [75].

Deep investigation tools These tools help users understand the scope and find the root cause of a threat. By choosing an entity of the interactive graph, users can ask questions and get to the root cause of that security threat [70].

Built-in queries A hunting tool which enables users to proactively hunt for security threats or anomalies in an organisation's data sources before an alert is triggered [70].

⁷Incidents are multiple related alerts collected together into a group

Notebooks Jupyter notebooks allow users to query, transform, analyse, and visualise Microsoft Sentinel data in a browser. Notebooks are used to extend the scope of what can do with your data, allowing features that aren't built into Microsoft Sentinel, such as Python machine learning features and customised data visualisation [70].

2.5.5 Content Hub

Content Hub is a marketplace for Microsoft Sentinel opened by Microsoft in 2021 and is meant to provide a unified method to access Microsoft Sentinel content [10]. In Content Hub, users are provided with a dashboard where they can find and install Out-of-the-box (OOTB) content and solutions in one step [76]. This means that Microsoft Sentinel provides users with pre-built templates, known as solutions, for standalone items or packaged solutions. The difference between standalone items and packaged solutions is that a standalone item only consists of one key feature provided in Microsoft Sentinel, whereas the packaged solution is a collection of one or more of the components from the Microsoft Sentinel key features, described in 2.5.4. These solutions can be downloaded from Content Hub, customised for a better fit to an environment's needs, and are often instantly usable [77]. The result is that organisations using Microsoft Sentinel can find and deploy solutions directly from Content Hub, instead of configuring their own data connectors, rules and workbooks.

To summarise, Content Hub is a marketplace in Microsoft Sentinel where different solutions are provided. These solutions consist of pre-built content packaged together, often containing data connectors, workbooks, analytics rules and playbooks.

In Content Hub you can find many solutions, but in our thesis, we will focus on the two following:

Microsoft Entra ID The solution consists of one data connector, two workbooks, 62 analytics rules, and 11 playbooks, and is provided by Microsoft. Microsoft Entra ID enables organisations to ingest different Microsoft Entra ID data logs using Diagnostic Settings into Microsoft Sentinel [78].

Microsoft Entra ID Protection The solution consists of one data connector, one analytics rule, and five playbooks, and is provided by Microsoft. Microsoft Entra ID Protection enables organisations to ingest Security alerts reported in Microsoft Entra ID Protection for risky users and events in Microsoft Entra ID [79].

2.5.6 Sentinel Dashboard

When using Sentinel you are provided with an Overview dashboard. This dashboard consists of widgets that represent the core components of Sentinel, for example, incidents, data, automation and analytics [80]. It is this dashboard or the

different widgets dashboard, you will use to gain information about the incidents in your system.

Incident widget In the incident widget of the dashboard you are provided with information about the incidents created in the last 24 hours. The information you are provided with includes the incident status, severity, closing classification, and creation time, and you are given a number for new, active and closed incidents [80].

Data widget In this widget you are presented with a graph containing information about different events, and their timestamps. You are also given information about your data connectors; how many of them are active, and how many of them are unhealthy [80]

Automation widget This widget gives you a summary of the automation rules activity. Here you get some information about closed incidents, time saved, and which actions are performed by the automation rules [80].

Analytics widget This widget gives you an overview of the status of the analytics rules you have in your environment [80].

2.6 False Results and Alert Fatigue

Alert fatigue is the concept that being constantly alert and overwhelmed by several incidents could lead security analysts to have poorer performance. This can be caused by several reasons, where the most important reason for this thesis is the number of *false positives* which show up [81].

A false positive is an alarm which points to an event which is not dangerous. A multitude of these alarms can lead to a security analyst potentially missing a true positive. This could result in a security breach [81]. We also have what is called a false negative, which is an error where a test is labelled as a negative when it should have been positive. In our context of cyber security, this would be the case of no detection for a successful attack [82].

To prevent alert fatigue Malwarebytes recommends changing to a detection and response security tool which can correlate events and remove benign incidents and false positives [81]. An example of such a tool is Sentinel as it has both SIEM and SOAR capabilities.

2.7 Attack Patterns

This section focuses on what identity-related attack patterns an attacker might use during an attack. The attack patterns described here are taken from the framework MITRE ATT&CK. The framework categorises attack patterns seen in real attacks to give organisations a better foundation for creating their cyber security defence

methodologies [83]. Presented below are the attack techniques which are related to user identities.

2.7.1 Reconnaissance

Reconnaissance is a tactic often performed at the beginning of an attack to gain information. This information is often then used to perform other tactics, often to achieve initial access to the target [84].

Gather Victim Identity Info (T1589)

A relevant technique for reconnaissance is *Gather Victim Identity Info* [33]. This technique is often used by attackers during the targeting stage of an attack. The information can be obtained in many ways; from scouring the internet to black markets. If an organisation has credentials on any such sites, users and applications with these credentials might be taken and used for initial access [33].

2.7.2 Initial Access

Gaining initial access is often the first step of an attack which happens on the target's systems. The goal of this tactic is to gain a foothold within the target to perform other attacks. MITRE has divided the act of gaining access to a system into two; one for the actual initial access and one for obtaining the credentials which could be used for initial access [85].

Valid Account (T1078)

A technique which is relevant to our case, with initial access to a cloud environment such as Azure, is *Valid account* [30]. An attacker might attempt to gain access to a system by using an already existing account, or by using the sign-in credentials. After signing in, the use of an existing account might help disguise malicious operations performed by the attacker [30]. Therefore, this technique is also associated with the defence evasion tactic [86].

2.7.3 Credential Access

Obtaining correct credentials for a system or user could open up the way for an attacker to gain initial access [87].

Multi-Factor Authentication Request Generation (T1621)

Gaining the correct sign-in credentials for a user might not be enough to gain access. If MFA is required for the sign-in, and the attacker does not have access to the MFA device, the attacker would need the user to accept the MFA request. A technique which attempts to achieve this is called *Multi-Factor Authentication*

Request Generation [88] and is often attempted by performing multiple sign-in attempts for the owner of the device to be bombarded with multiple MFA requests. This could lead to the user accepting one of the requests due to the user being fatigued by the amount of MFA requests, referred to as MFA fatigue. [88].

Steal Web Session Cookie (T1539)

Another technique for gaining credential access to internet services is through stealing another user's session cookie, or token. As cookies are both stored on the machine and sent together with other network traffic when requesting a site, an attacker has a multitude of other techniques which can be performed to gain a copy of the cookie. After gaining the copy of a cookie the attacker can use it to gain initial access to the related service [35].

These tokens are usually stored on the device and can be captured; either from looking through device memory or by eavesdropping on the browser traffic. One can usually find cloud sign-in cookies in the user browser under *content* or *application*. An adversary may attempt to download or copy over a token to another device [89].

In Azure the cookies *ESTAATH* and *ESTAATHPERSISTENT* are responsible for containing the session information [90]. One of these cookies alone is enough for authenticating a user and enough to bypass MFA in Azure during a *Steal Web Session Cookie* attack [91].

Brute Force (T1110)

The fourth technique for obtaining access credentials is through the use of "brute force" [92], which has several sub-techniques. This technique is used when an attacker does not have access to the credentials or has access to password hashes⁸ and is performed by systematically guessing different username and password combinations. The end goal of this technique is often for an attacker to gain access to a user to proceed with the *Valid Accounts* technique for initial access, but is also commonly used whenever an attacker has gained access to hashed passwords [92]. There are several sub-techniques, but only two of them are relevant to this thesis:

Password Guessing (T1110.001) This sub-technique is defined by an attacker having little to no knowledge about the password for a given account. The attack is performed by attempting different passwords for the account until the correct password is found. The attempts may be tailored to fit a password policy or be completely random. However, this method may be prone to failing due to locking out accounts as it attempts wrong passwords multiple

⁸A password hash is a user password which has been put through a hash function. These work one way and the result of each input string is unique. Storing passwords in this way is both considered safe and best practice [93].

times [94]. The default policy for *Azure Smart Lockout* in a public Azure tenant is to lock an account after the tenth failed attempt [25].

Password Spraying (T1110.003) This sub-technique employs a small sample of common passwords which are attempted towards a large sample of users. The reason for targeting a large sample of users is because account lockouts often happen when attempting a standard brute force attack with multiple password guesses towards a single user [32].

2.7.4 Privilege Escalation

In some cases, the user or method used by an attacker to get initial access might not have the permissions needed by the attacker to reach their objectives within the target network. This calls for the attacker to perform a technique which gives the attacker a higher level of privilege [95].

Valid Accounts (T1078)

As with initial access, a technique an attacker could use to achieve a higher level of permissions is to make use of an already valid account with the correct credentials [30].

Account manipulation (T1098)

Another technique an attacker might utilise is *Account manipulation* [31]. This technique is defined by changing settings for already existing users. Especially relevant for privilege escalation is the sub-technique *Account manipulation: Additional Cloud Roles* [96]. With this sub-technique, an attacker might add roles for a user they already have access to, in order to gain a user with the right privileges [96].

2.7.5 Command and Control

During an attack, the attacker needs to be able to communicate with and control the target. While employing this tactic, attackers usually attempt to disguise their traffic as being normal traffic or at least difficult to trace back to the attacker [97].

Proxy (T1090)

To prevent the system owner from identifying the attacker, an adversary might employ a technique called *Proxy* [34]. Here an attacker uses some sort of service to avoid a direct line of communication between the attacker and the target. Usually, this is combined with the use of an anonymisation service to mask the attacker's IP address. This technique makes it harder for the victim to find out who was behind the attack [34].

2.7.6 Gaining persistence

After getting the initial access an attacker might need to gain persistence in the environment [98].

Create account (T1136)

One of the methods which MITRE lists that focuses on identity is *create account*. With this technique the attacker achieves persistence through creating an account only they have access to [99].

Account Manipulation (T1098)

Account manipulation can also be used to achieve persistence in a system. This is usually through the sub-technique *Account manipulation: Additional Cloud Credentials* [100]. The goal of an attacker using this sub-technique is to alter an already existing cloud user by adding credentials. This will allow the attacker to continue to have access to the cloud user [100].

Chapter 3

Method

3.1 Outline of Chapter

This chapter describes the methodology employed to gather the necessary data to evaluate our research questions. As described in section 1.3 we have broken the main research question into four sub-questions. By answering each of these sub-questions we will have the foundation needed to answer the main research question. To answer the sub-questions we have created two cases based on sub-questions 1 and 2 (1.3).

3.1.1 Test Cases and Data Collecting

To collect the data needed to evaluate the first sub-research question, we will need to collect information from ID Protection and Sentinel. We will collect the number of detections performed by each system. In addition, we will inspect the information found about the alarms created in the ID Protection and Sentinel dashboards. This will provide us with a foundation that will make it possible to evaluate the first sub-research question.

The third sub-research question (1.3) provides us with the foundation for our test cases. We define two cases in which all attack techniques will be performed. The parameter will change for the two cases is *Sentinel, configured with the rulesets Microsoft presents in Content Hub for ME-ID and ID Protection, with best-practice CA enabled*. This will create two cases where one will have CA enabled and another will have CA disabled. These cases will help us more accurately discern what impact the use of CA has on the security setup's ability to detect attacks. Furthermore, we will be able to observe whether the use of CAPs will provide any further mitigating actions other than those performed by per-user MFA.

The answers to the two last sub-research questions (2 and 4) will be found by analysing the results given by the two test cases. To get an answer to sub-research question 2 of what alternations need to be done we will evaluate the results by checking the number of false positives and false negatives. We will also inspect whether there were any analytics rules or risk detections which could have been

expected to have triggered. This will make us able to evaluate what changes will be warranted to better the rulesets.

The final sub-research question will not directly help us evaluate the main research question. This question will rather help us evaluate the quality of our testing and if our tests will cover more ground than what Microsoft themselves provide as a guide for testing ID Protection (2.3.4).

3.2 Procedure for Testing

To test each of the cases we are going to attempt to simulate an attack which will use each of the relevant techniques presented by MITRE, described in 2.7. By performing these techniques ourselves we will be able to assess how the setup we are testing would hold up to a similar attack from an external source.

The goal of our thesis is to evaluate Microsoft's default recommended security detection systems. To achieve this, we will be performing a set of tests, detailed below. After each test has been performed we will be checking both the ID Protection dashboard and Sentinel incident dashboard for detections. As presented earlier, both ID Protection and Sentinel may require some time to generate an alarm. Because of this, we will wait 48 hours before being able to confidently assume that no alarm will show up¹.

To gather the results from our tests we will look for any alerts in the ID Protection and Sentinel dashboards as these are the platforms which are the focus of this thesis. If any alerts are found we will gather the information found there and evaluate it based on the criteria below. This information can be found in the Results chapter (4). Because ID Protection and Sentinel have different dashboards for their alerts, we will be able to see whether there are any discrepancies between the number of alerts detected by each service.

The results we are collecting from these tests will be evaluated based on the following criteria:

1. Was the attack successful?
2. Was the attack detected?
3. Was the attack mitigated?

The combination of each of these criteria will help us evaluate how well the recommended security setup from Microsoft is at detecting potential threats while using Sentinel and ID Protection. When evaluating whether the attack was *successful* or not, we will assess whether the attacker achieved their goal (gained access), or gained any meaningful information (such that they indeed have the correct credentials for the user).

We will consider the attack to have been *detected* if an alarm is created in the ID Protection dashboard or the Sentinel incident dashboard.

To determine whether the attack was *mitigated*, we will consider whether any mitigating actions were performed automatically by any service in the security

¹48 hours is the longest time interval before a risk detection shows up, 2.3.2.

setup which denies access. An example would be that during an attack we are prompted with MFA instead of being granted access. It is important to note that only the SIEM features of Sentinel are within the scope of our investigation due to our focus on attack detection. This means that the SOAR features are out of scope. Therefore, the question *was the attack mitigated?* does not apply to results gathered from Sentinel.

3.2.1 Testing Procedures

Valid Account

The first test we will conduct is for the technique *valid account* (described in 2.7.2). As described in 2.3.2 there are two risk detections which are designed to detect the use of this technique. Therefore, to simulate an attacker employing the *valid account* technique we will follow the guide for simulating *unfamiliar sign-in properties* as described in 2.3.4. Because the risk detection *atypical travel* also detects the use of a valid account from two distinct locations in a short period, we will also be testing this.

This technique undergoes a total of three tests: one in a new location, another using a new device, and a third involving atypical travel. Each test is going to follow the steps described in 2.3.4.

Success condition: For this attack to be considered successful the attacker is granted access. A partial success will be considered if MFA is requested.

MFA Request

To test the *MFA request* technique we will perform the steps outlined in 2.7.3. This test will be done in three different ways, each described in the following steps:

1. Navigate to `https://myapps.microsoft.com` and sign in using the sign-in credentials of the test user multiple times.
 2. Each time the MFA request arrives, turn it down.
 3. Accept the eleventh request.
 4. Await potential detection.
1. Navigate to `https://myapps.microsoft.com` and sign in using the sign-in credentials of the test user multiple times.
 2. Each time the MFA request arrives, wait till you get the notification "We didn't hear from you".
 3. Accept the eleventh request.
 4. Await potential detection.
1. Navigate to `https://portal.azure.com` and sign in using the sign-in credentials of the test user multiple times.
 2. Each time the MFA request arrives, wait till you get the notification "We didn't hear from you".
 3. Accept the eleventh request.

4. Await potential detection.

Success condition: For this attack to be considered successful the attacker needs to be granted access if the user accepts the MFA challenge. It will fail if the user is not able to accept the MFA challenge due to a number from the screen being required as input.

Gather Victim Info

Because Microsoft did not provide any method of simulating the *gather victim info* attack technique (2.7.1) for user identities, we have based this method on the *leaked credentials* steps outlined in 2.3.4. However, we are not testing on workload identities but rather on user identities. Therefore, we will upload the user credentials for a user on a public GitHub repository. The user will have no access to the cloud environment to safely perform this test. This is our procedure:

1. Using your GitHub account, create a public repository.
2. Create a file with the ".txt" extension and add the following [36]:

```
"AadUserPrincipleName : Victim1@ntnuinf t2504.onmicrosoft.com",  
"AadPassword : Toyota010234"
```

3. Commit the file and make sure to push the change to the repository.
4. Await potential detection after at least 8 hours.

Success condition: This attack does not have a success condition, because a potential attacker could find the credentials for the user as long as they are published, which is one of the steps during testing.

Proxy

The *proxy* technique (2.7.5) has a corresponding risk detection *anonymous IP address*, as outlined in 2.3.2. To test this technique we will perform the steps recommended by Microsoft in 2.3.4 for the risk detection *anonymous IP address*. This attack will be performed both on a user with MFA and a user without MFA. We are going to perform the attack with both these users to observe the difference between what is detected for a user with MFA enabled, and the MFA challenge is successful, and for a user with no MFA.

Success condition: For this attack to be considered successful the attacker needs to be granted access.

Brute Force

For us to test the *Brute Force* technique (2.7.3) we will be using a tool found online [101][102]. Due to ID Protection only creating an alarm if the correct credentials

are used for the sign-in attempt (2.3.2), we will be adding the correct username and password to the list of passwords which will be attempted².

We will attempt both password-guessing and password-spraying techniques using the procedure presented by the Atomic Red Team. We are using their procedure as they have more than 9'000 stars on the GitHub repository presenting their procedures³ and due to the procedure being directly linked to the specific MITRE techniques we are attempting to utilise [103].

Brute Force: Password Guessing To perform this test we will be using the procedure detailed by Atomic Red Team which targets a single user in ME-ID [101]. We will employ a short Powershell script (see A.1) which will take a list of passwords and a single username as input and attempts to perform the command **Connect-AzureAD** with each username-password combination [101]. The password list we will be using consists of 24 wrong passwords and a last successful one. The reason for using a list of 25 passwords is that a true password-guessing attack could reach high amounts of requests per second (2.7.3).

Brute Force: Password Spraying To perform this test we will also using a procedure detailed by the Atomic Red Team which targets multiple users in ME-ID [102]. We have a PowerShell script (see A.2) which will take as input a list of users and a single password. All the possible username-password combinations will then be attempted, while attempting to sign in using the same command as above, **Connect-AzureAD** [102]. We will be attempting this test using a list of 10 users.

Because all the users will have MFA enabled we know that the attempt will not provide a final successful sign-in attempt. However, an attacker utilising these techniques will know that they have used the correct credentials due to receiving a MFA challenge and not being refused outright.

Success condition: As mentioned in 2.7.3 the goal of a brute force attack is gaining access to credentials, therefore the attack is considered successful if the attacker can either sign in or use the correct credentials and be informed of needing MFA. If either of them occurs the attacker has been able to verify that they do have the correct access credentials and other attack techniques could be utilised further.

Steal Web Session Cookie

To perform the *Steal Web Session Cookie* we will create a procedure, based on 2.7.3, which allows a potential attacker to gain access to the user on another computer. This thesis does not focus on how to perform each attack, but rather on how our

²List of attempted username and password combinations can be found in the appendix together with the code.

³At the time of writing.

setup detects these attacks. Therefore, we are going to use a method of extracting the web session cookie which would not be available to most attackers.

As these cookies are stored on the host an attacker can gain these at any time after signing in, but in our case, we will capture them during the sign-in process (2.7.3). We will export the cookie using the cookie editor tool *Cookie Editor* [104]. The tool will be used as it has a large number of downloads and high reviews, as well as supporting functions for exporting and importing cookies. Leading to a fast method of exporting the cookie from one machine to another.

The cookie we will be targeting for the *Steal Web Session Cookie* technique is the *ESTSAUTH* cookie. The reason for targeting this cookie instead of the *ESTSAUTHPERSISTENT* cookie is that during the preliminary work for this thesis, while both learning the attack technique and designing these tests, we found no specific incentive for choosing the *ESTSAUTHPERSISTENT* cookie. Indeed, the same attack would be possible by using the *ESTSAUTHPERSISTENT* cookie, and an attack which extracts the persistent cookie, or both cookies, would be granted access to a session with a longer lifetime. However, we found that extraction of the *ESTSAUTH* cookie was easier than extraction of the *ESTSAUTHPERSISTENT* cookie while using our chosen tool. This was due to the *ESTSAUTH* cookie being configured earlier in the sign-in process enabling us to have a clearer and easier to understand testing procedure.

During testing we will perform the following steps, using two computers:

1. On computer 1, go to <https://login.microsoftonline.com/> and follow the normal sign-in procedure.
2. After accepting the MFA challenge, but before being forwarded to another domain, capture the cookies saved for the domain.
3. Send the cookie *ESTSAUTH* to computer 2.
4. On computer 2, go to <https://login.microsoftonline.com/>.
5. Use *Cookie Editor* to delete all previous cookies on computer 2.
6. Use *Cookie Editor* to import the *ESTSAUTH* cookie.
7. Refresh the page.

As mentioned previously in 2.7.3, the *ESTSAUTH* cookie alone holds all the information needed for signing in, including the MFA claim. Because an attacker then can bypass the MFA challenge, we will perform two different tests for this technique with the same procedure as above, but with two different procedures. In the first test, the two machines will be on the same network. During the second test, the two computers will be on two different networks with computer 2 on a previously unused IP address, this will be achieved with the use of a VPN. These are the parameters which change between the two tests:

1. Computer 2 has the same IP address which has been used previously.
2. Computer 2 has a different IP address which has not been used previously.

Success condition: This attack is considered successful if the attacker is granted access through the reuse of the cookie.

Account Manipulation

Account manipulation (2.7) is a technique used for persistence and privilege escalation, therefore we have two tests where the first focuses on privilege escalation and the last on persistence.

As mentioned in 2.7.4 we need a user with the proper privileges for manipulating an account and escalating privileges. For this, we will use a global admin user. Here is our procedure:

1. Sign-in with the admin account on <https://portal.azure.com>
2. Manoeuvre to Entra ID dashboard
3. Choose another user and change the user's role to global admin.

The second test has the goal of achieving persistence through utilising the *Additional Cloud Credentials* sub-technique (2.7.6). While performing this test we will be using a user which has been configured with MFA. We followed this procedure:

1. Sign-in with a user.
2. Go to <https://mysignins.microsoft.com/security-info>.
3. Register a new MFA device.
4. Sign in on another computer using the new MFA authentication method.

Success condition: Both of these attacks are considered successful if the attacker can change the settings for the user.

Create Account

To test the technique *create account* (2.7.6) we are going to use an admin account and create a new user. As mentioned in 2.7.6 we need a user with the proper privilege for creating an account. For this, we will use a user with global admin privileges. We are following this procedure:

1. Sign-in with the admin account on <https://portal.azure.com>.
2. Manoeuvre to Entra ID dashboard.
3. Create a new user, with the global admin role.

Success condition: This attack is considered successful if the attacker can create the user.

3.3 Entra ID Setup

3.3.1 Smart Lockout

In our setup, Smart Lockout will be in its default configuration, meaning it will lockout an account after 10 failed attempts, as mentioned in 2.2.4. This will only influence the *Brute Force: Password Guessing* test as it is the only test with enough failed attempts towards a user which would result in an account lockout.

3.3.2 Multi-Factor Authentication

For testing, the MFA challenge will be needed at each sign-in attempt. To achieve this, we activated *per-user MFA*⁴. We are initially planning to have per-user MFA enabled in our testing setup without the use of CA. However, following Microsoft's recommendations (2.4.3), we will disable per-user MFA when enabling CA [106].

3.3.3 Creating Test-Users

To set up all the users we need, we decided to do it by creating a script (*see A.1*) that made the users for us. We can also create each test user manually in Microsoft Entra, however, since we have 20 test users, creating a script that does this for us seems like the most efficient method. To do this we have to use Microsoft Graph PowerShell SDK⁵ to connect to our tenant. We can then automatically create all the users by using a CSV file (*see A.6*) with all the users, and then send them through a loop that creates the users and puts them into their respective groups.

Users

In our subscription, we have access to 25 Microsoft 365 E5 Developer licenses. Of these, only 20 can be used on our test users. To assign each licence to the users they need to have the property *Usage location* set to Norway. We set the correct property through scripting (A.1). Once this was done all licences could be assigned in ME-ID.

3.4 Microsoft Sentinel Deployment

To gain insight and gather results about how Microsoft Sentinel and Content Hub work, and how well they perform, we needed to deploy Microsoft Sentinel on our tenant and configure Content Hub with the desired solutions, which in our case is *Microsoft Entra ID* and *Microsoft Entra ID Protection*. Since the scope of our thesis focuses on *user identities*, see section 1.8, the two solutions we implemented in our environment will mainly have enabled the key features, see section 2.5.4, that are related specifically to user identities.

As described in section 1.3, the main goal of this thesis is to answer how well the rulesets provided by Microsoft in Microsoft Sentinel Content Hub (Content Hub) detect identity-based threats. To maximize the data collection from our Content Hub setup, we enabled as many analytics rules regarding user identities as possible. This approach will give us a better overview of how well each solution works so that we can further evaluate and discuss *how well these rulesets detect*

⁴According to Microsoft, per-user MFA is employed when users are required to use MFA every time they sign in [105].

⁵Microsoft Graph PowerShell SDK provides a way to use Microsoft Graph APIs with PowerShell. It makes all the APIs available in a simpler way and provides access to data stored across Microsoft 365 services [107].

threats. Also, as can be seen in section A.2, each analytics rule is very specific and will *not* cover a broad set of circumstances. Therefore, it is natural to enable as many rules as possible to ensure that we cover a wide range of potential attacks. This step is also crucial and necessary because, given the specificity of these rules to particular scenarios, it is unpredictable which analytics rule will be triggered by the attacks we simulate.

With all these analytics rules enabled, we will have the opportunity to collect all the necessary data to obtain a precise view of the system. This will further help us to answer our research questions (1.3) correctly.

3.4.1 Microsoft Entra ID Setup

As mentioned in section 2.5.5 the Microsoft Entra ID solution consists of one data connector, two workbooks, 62 analytics rules and 11 playbooks. The following table (3.1) will present which of the analytics rules we have enabled in our environment to detect attacks. For further information about what each of the specific analytics rules does, see A.2.

| Analytics rules: |
|-----------------------------------------------------------------|
| MFA Rejected by User |
| Attempt to bypass conditional access rule in Microsoft Entra ID |
| Failed login attempts to Azure Portal |
| Account Created and Deleted in Short TimeFrame |
| Account created or deleted by non-approved user |
| Attempts to sign in to disabled accounts |
| MFA Spamming followed by Successful login |
| Authentication Methods Changed for Privileged Account |
| Suspicious Sign In Followed by MFA Modification |
| Successful logon from IP and failure from a different IP |
| Distributed Password cracking attempts in Microsoft Entra ID |
| Password spray attack against Microsoft Entra ID application |
| Brute force attack against Azure Portal |
| Multiple admin membership removals from newly created admin |
| User Accounts - Sign in Failure due to CA Spikes |
| Privileged Accounts - Sign in Failure Spikes |
| Sign-ins from IPs that attempt sign-ins to disabled accounts |
| Explicit MFA Deny |
| New User Assigned to Privileged Role |
| User Assigned New Privileged Role |
| Bulk Changes to Privileged Account Permission |

Table 3.1: Analytics rules enabled in Microsoft Entra ID solution for the data connector Microsoft Entra ID

3.4.2 Microsoft Entra ID Protection Setup

As mentioned in section 2.5.5 the Microsoft Entra ID Protection solution consists of one data connector, one analytics rule and five playbooks. The following table (3.2) will present the analytics rule we have enabled in our environment to detect attacks. For further information about what the analytics rule do, see A.3.

| |
|--------------------------------------------------------------------|
| Analytics rules: |
| Correlate Unfamiliar sign-in properties and atypical travel alerts |

Table 3.2: Analytics rule enabled in Microsoft Entra ID Protection solution for the data connector Microsoft Entra ID Protection

3.4.3 How the Setup and Configuration is Done

As mentioned in section 2.5.5, Content Hub provides pre-built templates that you can enable and save without any further work. In our setup, we have followed these templates and deployed them as far as possible without any customisation to observe how well they naturally work.

However, some customisation and additional steps had to be done for the analytics rules to function. These steps will be explained in the following section.

Customisation of Analytics Rules

For most of the analytics rules, you could click right through the template and save them without any further work. However, for some of the rules, you had to enable the UEBA feature, as these analytics rules are based on behavioural profiles, making it possible to detect abnormalities.

3.4.4 Testing

For testing the solutions in Content Hub, we are going to simulate the attacks described in section 3.2.

The goal of this testing procedure is to determine the effectiveness of each solution in Content Hub for detecting attacks. Since we know which attacks are being simulated, we can easily measure what the solutions, or more specifically the analytics rules, discover and what they do not. This will serve as the metric we use to evaluate if Content Hub is suitable for organisations to use as an incident detection system, or if it does not perform as expected. By following this method, we can also identify unnecessary rules or any deficiencies in the solutions that may prevent the rules from being triggered under an attack.

These results will be collected from the *incident dashboard* inside of Sentinel. Here we can obtain an overview of each detected incident, including the analytics rule that triggered the detection and the time of detection. The timestamp of detection will also help us understand the specific simulated attack scenario that triggered the analytics rule.

When going into the incident dashboard in Sentinel we got an overview of all incidents detected based on our analytics rules. From here we clicked on the incident we wanted to examine, see figure 3.1.

| | | | | | | | |
|-------------------------------------|--------|----|----------------------------|---|----------------|--------------------|--------------------|
| <input checked="" type="checkbox"/> | Medium | 29 | Brute force attack agai... | 1 | Azure Sentinel | Microsoft Sentinel | 04/16/24, 04:11 PM |
| <input type="checkbox"/> | Medium | 28 | Password spray attack ... | 1 | Azure Sentinel | Microsoft Sentinel | 04/16/24, 04:05 PM |
| <input type="checkbox"/> | High | 27 | Authentication Metho... | 1 | Azure Sentinel | Microsoft Sentinel | 04/16/24, 03:42 PM |
| <input type="checkbox"/> | Low | 26 | Failed login attempts t... | 1 | Azure Sentinel | Microsoft Sentinel | 04/16/24, 02:36 PM |
| <input type="checkbox"/> | Medium | 25 | Password spray attack ... | 1 | Azure Sentinel | Microsoft Sentinel | 04/15/24, 04:05 PM |

Figure 3.1: Overview over the incident dashboard in Microsoft Sentinel

After this, we wanted the data collected when running the query the analytics rule is based on. To do this we clicked on *Events*, see figure 3.2.

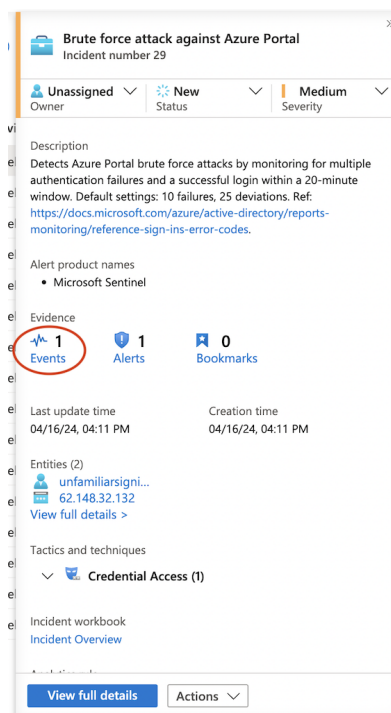
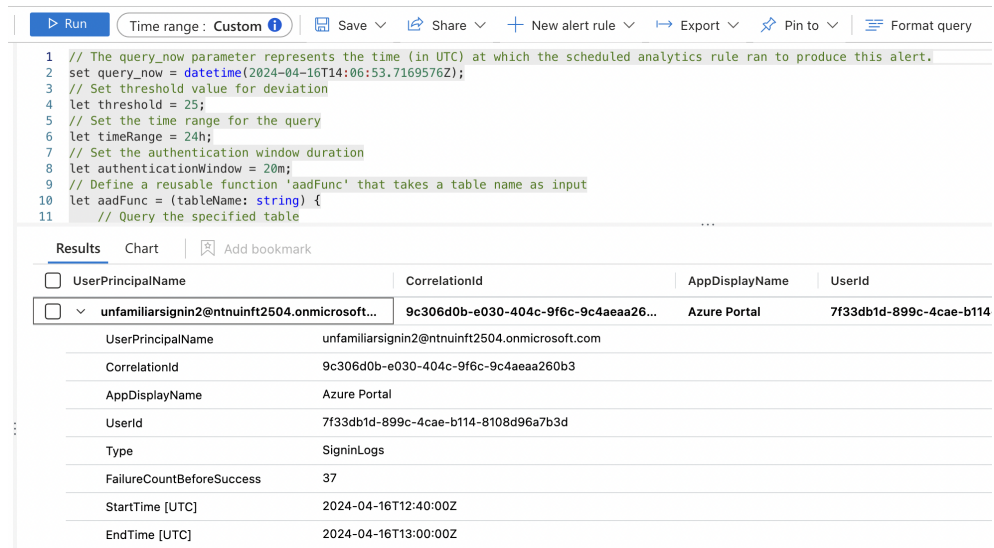


Figure 3.2: Overview over specific incident, with the *Events* marked

When doing this we got the query that ran to collect all the raw data, with the data placed in a clear table, see figure 3.3.



The screenshot shows a query editor interface with a 'Run' button and a 'Time range: Custom' dropdown. The query is as follows:

```

1 // The query_now parameter represents the time (in UTC) at which the scheduled analytics rule ran to produce this alert.
2 set query_now = datetime(2024-04-16T14:06:53.7169576Z);
3 // Set threshold value for deviation
4 let threshold = 25;
5 // Set the time range for the query
6 let timeRange = 24h;
7 // Set the authentication window duration
8 let authenticationWindow = 20m;
9 // Define a reusable function 'aadFunc' that takes a table name as input
10 let aadFunc = (tableName: string) {
11 // Query the specified table

```

The results table is displayed below the query, showing a single row of data for a failed sign-in attempt:

| UserPrincipalName | CorrelationId | AppDisplayName | UserId |
|-----------------------------------------------|-------------------------------------|----------------|--------------------------------------|
| unfamiliarsignin2@ntnuinf2504.onmicrosoft.com | 9c306d0b-e030-404c-9f6c-9c4aea260b3 | Azure Portal | 7f33db1d-899c-4cae-b114-8108d96a7b3d |

Below the table, a detailed view of the selected row is shown:

| | |
|---------------------------|-----------------------------------------------|
| UserPrincipalName | unfamiliarsignin2@ntnuinf2504.onmicrosoft.com |
| CorrelationId | 9c306d0b-e030-404c-9f6c-9c4aea260b3 |
| AppDisplayName | Azure Portal |
| UserId | 7f33db1d-899c-4cae-b114-8108d96a7b3d |
| Type | SigninLogs |
| FailureCountBeforeSuccess | 37 |
| StartTime [UTC] | 2024-04-16T12:40:00Z |
| EndTime [UTC] | 2024-04-16T13:00:00Z |

Figure 3.3: Overview over some of the query and table

3.5 Setup of Conditional Access Policies

This section outlines the approach taken to set up and configure CAPs based on the Zero Trust template from Microsoft. The setup followed standard practices and mainly used Microsoft Entra to deploy CAPs.

3.5.1 Tools used

To get the best possible setup of CAPs for our testing scenarios, we chose to use the 14 policies from the template based on Zero Trust (see 2.4.3 and figure 3.4). By utilising this method, multiple CAPs will be set up without the need for manual configuration for each policy.

| Policy name | State |
|--------------------------------------------------------------------------------|-------------|
| Block access for unknown or unsupported device platform | Report-only |
| Block access for users with Insider Risk (Preview) | Report-only |
| Block legacy authentication | Report-only |
| No persistent browser session | Report-only |
| Require approved client apps or app protection policies | Report-only |
| Require compliant or hybrid Azure AD joined device or multifactor authentic... | Report-only |
| Require multifactor authentication for Azure management | Report-only |
| Require multifactor authentication for Microsoft admin portals | Report-only |
| Require multifactor authentication for admins | Report-only |
| Require multifactor authentication for all users | Report-only |
| Require multifactor authentication for guest access | Report-only |
| Require multifactor authentication for risky sign-ins | Report-only |
| Require password change for high-risk users | Report-only |
| Securing security info registration | Report-only |

Figure 3.4: Which Conditional Access policies are configured from the Zero-Trust template, see section 2.4.3

3.5.2 Setup and Configuration

When setting up the policies, we are following Microsoft's recommendations [108] to create them in report-only mode, see figure 3.4. In this mode, policies are evaluated but not enforced during sign-in. The results are logged in both the CA and Report-only tabs of the Sign-in log details [42].

As one of the goals of the thesis is performing AB-tests, starting with the policies in report-only will make the testing process much easier. This approach eliminates the need to manually disable all CAPs when conducting AB-tests (see 3.1.1 and 3.2) without CA. However, they will be enabled during the second test case when testing with CA, to achieve a more accurate understanding of which policies are triggered and how they affect detection in Sentinel and ID Protection.

3.6 Evaluation of The Chosen Method

Every method has some positives and a few drawbacks. In this section, we will evaluate our chosen method. This involves the method used for performing our tests, how we collected and evaluated our data, and how our setup for ME-ID and Sentinel might have impacted our results.

3.6.1 Performing The Attack Techniques

Most attackers are likely to combine different techniques during an attack, however, we chose to evaluate each technique individually. We determined that a method of combining different techniques would create a large number of possible tests from each possible combination which could make it difficult to discern which technique triggered a given response. Because of these two potential difficulties, we decided that it would be better to test each technique individually. However, while testing the *steal web session cookie* we did perform a combination of different techniques. The reasoning behind this was that this attack technique could be more dangerous as it can bypass MFA. This focus allowed us to better assess our setup's ability to detect the attack techniques.

During testing, we often reused the machines and IP addresses which we used for signing in. The reason for us reusing these parameters was that they were not relevant for many of the tests we performed. Therefore we saw no benefit in creating a new VM or using a new VPN for each of the tests, and only changed these for the test where this was specified. However, due to ID Protection using machine learning in their detection algorithm, some of our attacks may be detected and labelled as false positives, without providing us with any information about the reasoning of their decision.

In many ways, machine learning is a black box. Microsoft provides little insight into what makes their detection algorithm for ID Protection mark an event as a false positive. This is especially relevant for the techniques which have corresponding risk detections in ID Protection (2.3.2). As mentioned in 2.3.4, Microsoft says that only a few of the risk detections can be triggered manually. The result of this is that we are unable to trust a result of *no detection* in ID Protection on these techniques. To mitigate this problem we will consider both a case of *detection* and *no detection* for results from ID Protection for the techniques which have correlating risk detections (2.3.2).

3.6.2 Entra Setup

Each user in our tenant has been set up with MFA. This setup ensures that our results are consistent and reliable across all test users, as some of the tests require MFA to be enabled to function correctly.

An advantage of our approach and setup is the minimal presence of unnecessary elements that could potentially distort our results. By maintaining a simplified setup with few other influencing factors, we can clearly show the outcomes derived from our testing.

However, there are inherent limitations due to the non-specific nature of our approach, which does not adapt to the specific circumstances of a particular business. By adhering strictly to Microsoft's recommendations, our CAP setup might not address certain practical considerations that a typical business would face. For instance, our CA setup does not specifically address geographical login restrictions, which are often critical for organisations without international presence or

those not requiring travel abroad. This oversight might result in a CAP setup that, while theoretically robust, could miss practical safeguards important in real-world business scenarios.

Moreover, the reliance on predefined threat models in the Zero Trust template could potentially have implications for our testing process. Most of the CAPs are policies that require MFA, which may interfere with several of our planned test scenarios, particularly those involving rapid or automated access, thus potentially hindering effective data collection efforts.

3.6.3 Sentinel Setup

When we configured Sentinel, our goal was to gain as much information as possible about the different analytics rules' ability to detect different attacks. To reach this goal we enabled every analytics rule that was based on *user identities* found in the ME-ID and ID Protection solutions from Content Hub.

The advantage of this setup is that it will give the most amount of coverage for the different attack scenarios which might arise, regarding user identities. In practice, Sentinel should have detected the different attack techniques we tested, provided there was a suitable analytics rule. Only enabling certain analytics rules would have led to a possibly incomplete view of the chosen solutions' ability to detect our attacks.

However, there is a disadvantage to this chosen method; it is possible to receive multiple alarms in your system, which may alert you about the same attack. This means that you may receive more alerts than necessary, which is not always beneficial due to the possibility of alert fatigue, as mentioned in 2.6.

3.6.4 Summary

As a final evaluation of our method, we find the method to be an effective method for being able to evaluate what attack techniques are going to be detected by the recommended security setup. However, our setup is not equal to what a real organisation would set up. This is with reason, as we are not trying to copy a real organisation. Our method will provide answers that an organisation could use to help evaluate how they could fit the recommended setup for CAPs and Sentinel rules from Content Hub for ME-ID and ID Protection.

Chapter 4

Results

4.1 Outline of Chapter

This chapter will present our results, aiming to provide an overview of the outcomes achieved using each system. We strive to gather enough insights from these systems to effectively address the research questions described in section 1.3.

To ensure clarity throughout this chapter, we have chosen to discuss each attack separately and present the findings from each system individually about that attack. The chapter is structured with 15 attacks, followed by descriptions of the results obtained from Sentinel, ID Protection, CA, and, where applicable, any additional descriptions involving when multiple systems were configured simultaneously.

Bear in mind, testing criteria 3, *Was the attack mitigated?*, is not applicable in the Sentinel results (3.2).

4.2 Valid Account

4.2.1 Valid Account: New location

Results from Sentinel

When we carried out this attack, described in 2.3.4, Sentinel created *one* incident triggered by the analytics rule *MFA explicitly deny*. The result for the incident is described in table 4.1 and is provided below¹.

¹The table describing the analytics rule *MFA explicit deny* is comprehensive. A full version of this table can be found in appendix A.4.

| | |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TimeGenerated [UTC] | 2024-04-17T10:44:10.0055312Z |
| OperationName | Sign-in activity |
| Category | SignInLogs |
| ResultType | 500121 |
| ResultSignature | None |
| ResultDescription | Authentication failed during strong authentication request. |
| DurationMs | 0 |
| Resource | Microsoft.aadiam |
| Identity | Lea |
| Level | 4 |
| Location | IN |
| AppDisplayName | My Apps |
| AuthenticationContextClassReferences | id:"urn:user:registersecurityinfo",detail:"previouslySatisfied" |
| AuthenticationDetails | Previously satisfied, now; Authentication in progress" |
| AuthenticationRequirement | multiFactorAuthentication |
| AuthenticationRequirementPolicies | ["requirementProvider":"user",detail:"Per-user MFA"] |
| ClientAppUsed | Browser |
| ConditionalAccessStatus | notApplied |
| CreatedDateTime [UTC] | 2024-04-17T10:41:23.5448174Z |
| IpAddress | 45.137.126.165 |
| LocationDetails_dynamic | "city":"Chennai","state":"Tamil Nadu","countryOrRegion":"IN" |
| MfaDetail_dynamic | "authMethod":"Mobile app notification" |
| ProcessingTimeInMilliseconds | 95 |
| RiskDetail | none |
| RiskLevelAggregated | none |
| RiskLevelDuringSignIn | none |
| ResourceDisplayName | Microsoft Graph |
| Status_dynamic | "errorCode":500121,"additionalDetails":"MFA denied; user declined the authentication","failureReason":"Authentication failed during strong authentication request." |
| TokenIssuerType | AzureAD |
| UserDisplayName | Lea |
| UserPrincipalName | lea@ntnuinf2504.onmicrosoft.com |
| AADTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| UserType | Member |
| ResourceTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| HomeTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| AutonomousSystemNumber | 62240 |
| AuthenticationProtocol | none |
| CrossTenantAccessType | none |
| Type | SignInLogs |
| PublicIP | 45.137.126.165 |
| Name | lea |
| UPNSuffix | ntnuinf2504.onmicrosoft.com |

Table 4.1: Analytics Rule: MFA explicitly deny

Due to the fulfilment of the partial success condition (3.2.1), this attack can be seen as partially successful.

This attack can be summarised:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

The user showed up in ID Protection Risk Detections with the detection timing set as *Real-time* in ID Protection. The detection type *Unfamiliar sign-in properties* was detected (see figure 4.1).

Due to the fulfilment of the partial success condition (3.2.1), this attack can be seen as partially successful.

The attack was mitigated, in the form of an account not gaining access to the environment due to MFA being prompted. This mitigation came from the *per-user MFA*.

The results from ID Protection can be summarised:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Yes*

| User detections | | | | | | |
|------------------------|---------|---------------|-------------------------|---------------------------|---------------|---------------|
| Detection time ↑↓ | User ↑↓ | IP address ↑↓ | Location | Detection type ↑↓ | Risk state ↑↓ | Risk level ↑↓ |
| 4/17/2024, 12:22:40 PM | Lea | 45.137.126.8 | Chennai, Tamil Nadu, IN | Unfamiliar sign-in pro... | At risk | High |

Figure 4.1: From ID Protection risk detections: Risk detection registered

| Risk Detection Details | |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| User's risk report User's sign-ins User's risky sign-ins Linked risky sign-in User's risk detections | |
| Detection type | Unfamiliar sign-in properties |
| Risk state | At risk |
| Risk level | High |
| Risk detail | - |
| Attack type(s) | Access using a valid account (Detected at Sign-In) |
| Source | Identity Protection |
| Detection timing | Real-time |
| Activity | Sign-in |
| Detection time | 4/17/2024, 12:22 PM |
| Detection last updated | 4/17/2024, 12:38 PM |
| Token issuer type | Microsoft Entra ID |
| Additional info | The following properties of this sign-in are unfamiliar for the given user: ASN, browser, device, IP, location, Exchange Active Sync ID, tenant IP subnet. |
| Sign-in time | 4/17/2024, 12:22 PM |
| IP address | 45.137.126.8 |
| Sign-in location | Chennai, Tamil Nadu, IN |
| Sign-in client | Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7) AppleWebKit/605.1.15 (KHTML, like Gecko) Version/17.1 Safari/605.1.15 |

Figure 4.2: From ID Protection risk detections: Details of triggered risk detection.

Results with Conditional Access Enabled

After the attacker failed the MFA challenge, four CAPs were applied, as seen in figure 4.5:

- *Require multifactor authentication for admins* - Status: Failure
- *Require multifactor authentication for all users* - Status: Failure
- *No persistent browser session* - Status: Success
- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users* - Status: Failure

The policies succeeded in blocking the user from signing in, thereby they mitigated the attack, as the attacker did not succeed the MFA challenge (see figure 4.4). This resulted in the policy results being a failure, as seen in figure 4.5. The attacker was prompted with MFA thereby making this attack a partial success (3.2.1).

The attack was detected in ID Protection with the identical risk detection, thereby no new risk detection was triggered with CA enabled. However, Sentinel created two new incidents. This attack can be summarised:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Yes*

| | | | | | | | |
|-----------------------|-----------------------|-----|------------|---------|---------------|--------------------------|---------|
| 4/19/2024, 1:25:32 PM | fd9584fb-b2c3-4eba... | Lea | OfficeHome | Failure | 45.130.203.89 | Al Qahirah, Al Qahira... | Failure |
|-----------------------|-----------------------|-----|------------|---------|---------------|--------------------------|---------|

Figure 4.3: From ME-ID sign-in logs: User trying to sign in from Egypt, but failing MFA challenge

Activity Details: Sign-ins ✕

Basic info
Location
Device info
Authentication Details
Conditional Access
Report-only
⋮

| | |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Date | 4/19/2024, 1:25:32 PM |
| Request ID | fd9584fb-b2c3-4eba-a6d3-23c955094b00 |
| Correlation ID | 7bc34ea-e5ff-42bc-9e04-4aec61c2c2e |
| Authentication requirement | Multifactor authentication |
| Status | Failure |
| Continuous access evaluation | No |
| Sign-in error code | 500121 |
| Failure reason | Authentication failed during strong authentication request. |
| Additional Details | The user didn't complete the MFA prompt. They may have decided not to authenticate, timed out while doing other work, or has an issue with their authentication setup. |
| Troubleshoot Event | <p>Follow these steps:</p> <p>Launch the Sign-in Diagnostic.</p> <ol style="list-style-type: none"> 1. Review the diagnosis and act on suggested fixes. |
| User | Lea |
| Username | lea@ntnuinf2504.onmicrosoft.com |
| User ID | 90f53014-b7e3-4dfe-b91c-6927384f673c |

Figure 4.4: From ME-ID sign-in logs: Activity details sign-ins, user is unable to sign in due to failed MFA challenge

Activity Details: Sign-ins ✕

Basic info
Location
Device info
Authentication Details
Conditional Access
Report-only
⋮

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ | |
|------------------------------------------------------------------|------------------------------|-------------------------------|-----------|---|
| Require multifactor authentication for admins | Require multifactor authe... | | Failure | ⋮ |
| Require multifactor authentication for all users | Require multifactor authe... | | Failure | ⋮ |
| No persistent browser session | | Sign-in frequency, Persist... | Success | ⋮ |
| Require compliant or hybrid Azure AD joined... | Require multifactor authe... | | Failure | ⋮ |

Figure 4.5: From ME-ID sign-in logs: Activity details, CAPs applied

Sentinel After Enabling CA

After we turned on CA we got different incidents in Sentinel, even though we followed the same approach for the attack.

We got *two* new incidents in Sentinel, triggered by the analytics rules *Successful logon from IP and failure from a different IP* and *Correlate Unfamiliar sign-in prop-*

erties & atypical travel alerts. The results for each of these incidents are presented in table 4.2 and table 4.3.

| | |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| UserPrincipalName | lea@ntnuinft2504.onmicrosoft.com |
| SuccessIPAddress | 45.130.203.89 |
| SuccessLocation | EG |
| AppDisplayName | Azure Portal |
| FailedIPAddress | 62.148.32.132 |
| FailedLocation | NO |
| ResultType | 50097 |
| ResultDescription | Device Authentication Required - DeviceId -DeviceAltSecId claims are null OR no device corresponding to the device identifier exists. |
| Type | SigninLogs |
| FailedLogonTime [UTC] | 2024-04-19T11:27:36.8125664Z |
| SuccessLogonTime [UTC] | 2024-04-19T11:25:47.9508079Z |
| timestamp [UTC] | 2024-04-19T11:25:47.9508079Z |
| Name | lea |
| UPNSuffix | ntnuinft2504.onmicrosoft.com |
| AccountUPN | lea@ntnuinft2504.onmicrosoft.com |
| GroupMembership | NTNU |
| AssignedRoles | GlobalAdministrator |
| UserType | Member |
| UserAccountControl | Member |
| UserInsights | "AccountDisplayName":"Rest 9", "AccountObjectID":"472b12ff-f481-4139-a00c-4473e3a9309b" User |
| | "AccountDisplayName":"Lea", "AccountObjectID":"90f53014-b7e3-4dfe-b91c-6927384f673c" |
| | "AccountDisplayName":"Amund", "AccountObjectID":"28449167-a66b-49b5-ba28-1382de195b70" |
| | "AccountDisplayName":"Rest 2", "AccountObjectID":"8f48cf18-9863-4827-aa9e-37c1c39f02a2" User |
| | "AccountDisplayName":"Rest 1", "AccountObjectID":"b078247b-4264-4a30-bfed-99bfa1a0d351" User |
| | "AccountDisplayName":"Travel 2", "AccountObjectID":"83d4f0b4-0c5f-4798-a545-13dcbe3d98de" User |
| | "AccountDisplayName":"Amund Strømsnes", "AccountObjectID":"d6fbb28a-4b6d-4ae4-b43f-626072d66b2c" Fredrik |
| DeviceInsights | "UserAgentFamily":"Firefox" |
| | "UserAgentFamily":"Chrome" |
| | "UserAgentFamily":"Edge" |
| IPInvestigationPriority | 38 |
| UEBARiskScore | 38 |

Table 4.2: Analytics Rule: Successful logon from IP and failure from a different IP

| | |
|----------------------------------------|----------------------------------|
| UserAccount | lea@ntnuinft2504.onmicrosoft.com |
| Alert_UnfamiliarSignInProps_Name | Unfamiliar sign-in properties |
| Alert_UnfamiliarSignInProps_Severity | High |
| Alert_UnfamiliarSignInProps_Time [UTC] | 2024-04-19T20:15:39.6669795Z |
| Alert_AtypicalTravels_Name | Atypical travel |
| Alert_AtypicalTravels_Severity | Low |
| Alert_AtypicalTravels_Time [UTC] | 2024-04-19T20:16:15.887109Z |
| TimeDelta | -00:00:36.2201295 |
| CurrentLocation | Al Qahirah, Al Qahirah, EG |
| PreviousLocation | NO |
| CurrentIPAddress | 45.130.203.89 |
| PreviousIPAddress | 62.148.32.132 |
| UserName | Lea |
| UserEmailName | lea |
| UPNSuffix | ntnuinft2504.onmicrosoft.com |

Table 4.3: Analytics Rule: *Correlate Unfamiliar sign-in properties & atypical travel alerts*

Due to the fulfilment of the partial success condition, described in 3.2.1, this attack can be seen as partially successful.

This attack can be summarised:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Not applicable*

4.2.2 Valid Account: New device

Results from Sentinel

When we carried out this attack, described in 2.3.4, Sentinel created two incidents triggered by the analytics rules *MFA explicitly deny* and *Failed login attempts to the Azure Portal*. The result for the incidents is described in table 4.4 and table 4.5 and is provided below.

| | |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TimeGenerated [UTC] | 2024-04-18T08:42:02.8376811Z |
| OperationName | Sign-in activity |
| Category | SignInLogs |
| ResultType | 500121 |
| ResultSignature | None |
| ResultDescription | Authentication failed during strong authentication request. |
| DurationMs | 0 |
| Resource | Microsoft.aadiam |
| Identity | Amund |
| Level | 4 |
| Location | NO |
| AppDisplayName | My Apps |
| AuthenticationContextClassReferences | id:"urn:user:registersecurityinfo",detail:"previouslySatisfied" |
| AuthenticationDetails | Previously satisfied, now; MFA denied; user declined the authentication |
| AuthenticationRequirement | multiFactorAuthentication |
| AuthenticationRequirementPolicies | ["requirementProvider":"user",detail:"Per-user MFA"] |
| ClientAppUsed | Browser |
| ConditionalAccessStatus | notApplied |
| CreatedDateTime [UTC] | 2024-04-18T08:29:22.9376554Z |
| IpAddress | 62.148.32.132 |
| LocationDetails_dynamic | "city":"Fornebu","state":"Akershus","countryOrRegion":"NO" |
| MfaDetail_dynamic | "authMethod":"Mobile app notification" |
| ProcessingTimeInMilliseconds | 119 |
| RiskDetail | none |
| RiskLevelAggregated | none |
| RiskLevelDuringSignIn | none |
| ResourceDisplayName | Microsoft Graph |
| Status_dynamic | "errorCode":500121,"additionalDetails":"MFA denied; user declined the authentication","failureReason":"Authentication failed during strong authentication request." |
| TokenIssuerType | AzureAD |
| UserDisplayName | Amund |
| UserPrincipalName | amund@ntnuinf2504.onmicrosoft.com |
| AADTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| UserType | Member |
| ResourceTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| HomeTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| AutonomousSystemNumber | 13243 |
| AuthenticationProtocol | none |
| CrossTenantAccessType | none |
| Type | SignInLogs |
| PublicIP | 62.148.32.132 |
| Name | amund |
| UPNSuffix | ntnuinf2504.onmicrosoft.com |

Table 4.4: Analytics Rule: MFA explicitly deny

| | |
|-------------------|---------------------------------------------------------------|
| UserPrincipalName | amund@ntnuinft2504.onmicrosoft.com |
| UserId | 28449167-a66b-49b5-ba28-1382de195b70 |
| UserDisplayName | Amund |
| Status | 50079: User needs to enroll for second factor authentication. |
| FailedLogonCount | 36 |
| IPAddress | 62.148.32.132 |
| IPAddressCount | 1 |
| AppDisplayName | Azure Portal |
| Browser | Firefox 115.0 |
| OS | Linux |
| FullLocation | NO Akershus Fornebu |
| Type | AADNonInteractiveUserSignInLogs |
| StartTime [UTC] | 2024-04-18T08:08:35.890972Z |
| EndTime [UTC] | 2024-04-18T08:10:09.0017709Z |
| Name | amund |
| UPNSuffix | ntnuinft2504.onmicrosoft.com |

Table 4.5: Analytics Rule: *Failed login attempts to the Azure Portal*

Since the attacker got MFA requested this attack can be considered a partial success, due to the fulfilment of the partial success condition outlined in 3.2.1.

This can be summarised as:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

The test did not trigger any risk detections in ID Protection, meaning it was not detected.

According to the success condition outlined in 3.2.1, this attack can be viewed as a partial success.

The attack can be seen as mitigated since the attacker did not gain access to the system, due to being prompted with MFA.

The results from this test can be summarised:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Yes*

Results with Conditional Access Enabled

After the attacker failed the MFA challenge, four CAPs were applied (see figure 4.6):

- *Require multifactor authentication for admins* - Status: Failure
- *Require multifactor authentication for all users* - Status: Failure
- *No persistent browser session* Status: Success
- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users* - Status: Failure

The policies succeeded in blocking the admin user from signing in, therefore the attack was mitigated. The failed policies can be seen in figure 4.6. This attack is a partial success based on the success condition outlined in 3.2.1.

The attack was detected in Sentinel with the same incident being created. No risk detections were triggered.

This can be summarised:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Yes*

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details Conditional Access Report-only ...

🔍 Search

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ |
|------------------------------------------------------------------|------------------------------------|-----------------------------|-----------|
| Require multifactor authentication for admins | Require multifactor authentication | | Failure |
| Require multifactor authentication for all users | Require multifactor authentication | | Failure |
| No persistent browser session | | Sign-in frequency, Persi... | Success |
| Require compliant or hybrid Azure AD joined ... | Require multifactor authentication | | Failure |

Figure 4.6: From ME-ID Sign-in logs: The four CAPs that were applied

4.2.3 Valid Account: Atypical Travel

Results from Sentinel

This attack was not detected by Sentinel.

Due to the fulfilment of the partial success condition, outlined in 3.2.1, this attack can be seen as partially successful.

This attack can be summarised:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

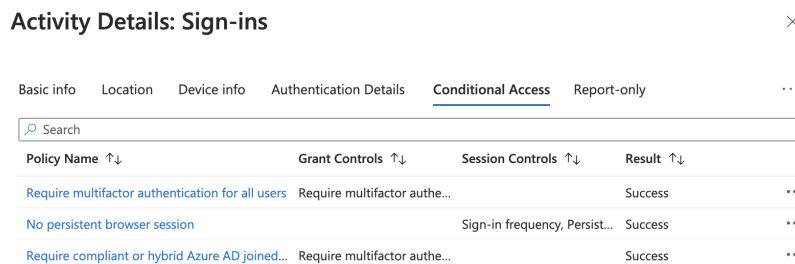
The test did not create any risk detection in ID Protection. Because MFA was requested, this test can be considered a partial success based on the success condition outlined in 3.2.1. The attack was not mitigated as the MFA challenge was completed. This test can be summarised:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

Results with Conditional Access Enabled

These are the three CAPs that were applied and their result status (see figure 4.7):

- *Require multifactor authentication for all users - Status: Success*
- *No persistent browser session - Status: Success*
- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users - Status: Success*



Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details **Conditional Access** Report-only ...

Search

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ | |
|--------------------------------------------------|------------------------------------|-------------------------------|-----------|-----|
| Require multifactor authentication for all users | Require multifactor authentication | | Success | ... |
| No persistent browser session | | Sign-in frequency, Persist... | Success | ... |
| Require compliant or hybrid Azure AD joined... | Require multifactor authentication | | Success | ... |

Figure 4.7: From ME-ID Sign-in logs: The three CAPs that were applied

Due to the condition of MFA being met, the CAPs resulted in success, thereby giving the attacker access to sign in and the attack was not mitigated. This test can be seen as a partial success, having met the success condition outlined in 3.2.1.

The test did not create any incidents in Sentinel and no risk detections in ID Protection with CA enabled.

This can be summarised:

1. Was the attack successful? *Yes, partially*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

4.3 MFA Request

4.3.1 MFA Request (explicit deny)

Results from Sentinel

When we carried out this attack, described in 3.2.1, Sentinel created *two* incidents triggered by the analytics rules *MFA explicitly deny* and also *MFA Spamming followed by Successful login*. The results for the incidents are presented in table 4.6 and table 4.7 and are provided below.

| | |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TimeGenerated [UTC] | 2024-04-09T07:51:43.4558273Z |
| OperationName | Sign-in activity |
| Category | SignInLogs |
| ResultType | 500121 |
| ResultSignature | None |
| ResultDescription | Authentication failed during strong authentication request. |
| DurationMs | 0 |
| Resource | Microsoft.aadiam |
| Identity | Rest User 3 |
| Level | 4 |
| Location | NO |
| AppDisplayName | My Apps |
| AuthenticationContextClassReferences | id:"urn:user:registersecurityinfo",detail:"previouslySatisfied" |
| AuthenticationDetails | Previously satisfied, now; MFA denied, user declined authentication |
| AuthenticationRequirement | multiFactorAuthentication |
| AuthenticationRequirementPolicies | ["requirementProvider":"user",detail:"Per-user MFA"] |
| ClientAppUsed | Browser |
| ConditionalAccessStatus | notApplied |
| CreatedDateTime [UTC] | 2024-04-09T07:45:21.0629194Z |
| IpAddress | 62.148.32.132 |
| LocationDetails_dynamic | "city":"Fornebu","state":"Akershus","countryOrRegion":"NO" |
| MfaDetail_dynamic | "authMethod":"Mobile app notification" |
| RiskDetail | none |
| RiskLevelAggregated | none |
| RiskLevelDuringSignIn | none |
| ResourceDisplayName | Microsoft Graph |
| Status_dynamic | "errorCode":500121,"additionalDetails":"MFA denied; user declined the authentication","failureReason":"Authentication failed during strong authentication request." |
| TokenIssuerType | AzureAD |
| UserDisplayName | Rest User 3 |
| UserPrincipalName | restuser3@ntnuinf2504.onmicrosoft.com |
| UserType | Member |
| ResourceTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| HomeTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| Type | SignInLogs |
| PublicIP | 62.148.32.132 |
| Name | restuser3 |
| UPNSuffix | ntnuinf2504.onmicrosoft.com |

Table 4.6: Analytics Rule: MFA explicitly deny

| | |
|----------------------|---------------------------------------|
| UserPrincipalName | anonuser2@ntnuinf2504.onmicrosoft.com |
| IPAddress | 62.148.32.132 |
| State | Akershus |
| City | Fornebu |
| Region | NO |
| FailedAttempts | 18 |
| SuccessfulAttempts | 1 |
| InvolvedOS | ["MacOS"] |
| InvolvedBrowser | ["Safari 16.4"] |
| StartTime [UTC] | 2024-04-19T08:23:21.2005118Z |
| EndTime [UTC] | 2024-04-19T08:27:42.1646911Z |
| AuthenticationWindow | 00:04:20.9641793 |
| Name | anonuser2 |
| UPNSuffix | ntnuinf2504.onmicrosoft.com |

Table 4.7: Analytics Rule: *MFA Spamming followed by Successful login*

Since the user had to write in a number to accept the MFA challenge this attack is considered to have failed, as described in 3.2.1.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

This attack was not detected in ID Protection.

The success condition for 3.2.1 was not met and the attack was therefore not successful.

Due to having to write in the right number from the MFA challenge, this attack can be considered to have been mitigated. This mitigation comes from *per-user MFA*.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Yes*

Results with Conditional Access Enabled

For each time the MFA request was turned down these CAPs failed due to the conditions of MFA not being met (see figure 4.8):

- *Require multifactor authentication for all users* - Status: Failure
- *No persistent browser session* - Status: Success

- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users* - Status: Failure

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details Conditional Access Report-only ⋮

🔍 Search

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ | |
|------------------------------------------------------------------|---------------------------------------|----------------------------------|-----------|---|
| Require multifactor authentication for all users | Require multifactor authentication... | | Failure | ⋮ |
| No persistent browser session | | Sign-in frequency, Persistent... | Success | ⋮ |
| Require compliant or hybrid Azure AD joined... | Require multifactor authentication... | | Failure | ⋮ |

Figure 4.8: From ME-ID Sign-in logs: All the CAPs that were applied after denying the MFA challenge

Once the MFA request was completed the CAPs that first failed were now successful, due to the condition of MFA being met (4.9)

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details Conditional Access Report-only ⋮

🔍 Search

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ | |
|-------------------------------------------------------|---------------------------------------|------------------------------------|-----------|---|
| Require multifactor authentication... | Require multifactor authentication... | | Success | ⋮ |
| No persistent browser session | | Sign-in frequency, Persistent b... | Success | ⋮ |
| Require compliant or hybrid A... | Require multifactor authentication... | | Success | ⋮ |

Figure 4.9: From ME-ID Sign-in logs: All the CAPs that were applied after completing the MFA challenge

Because MFA was requested this attack was not successful based on the success condition outlined in 3.2.1. The attack was mitigated due to MFA being requested stopping the attacker from getting access.

The attack was detected in Sentinel however, after we enabled CA no new incident in Sentinel or any risk detections in ID Protection was triggered.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Yes*

4.3.2 MFA Request (no answer)

Results from Sentinel

After we carried out this attack, described in 3.2.1, against `myapps.microsoft.com`, no detection was done by Sentinel.

Since the user had to write in a number to accept the MFA challenge, this attack is considered a failure, as described by the success condition in 3.2.1.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Not applicable*

However, when we carried out this attack, described in 3.2.1, against `portal.azure.com`, Sentinel created one incident triggered by the analytics rule *Brute Force attack against Azure Portal*. The result for the incident is described in table 4.8 provided below.

| | |
|---------------------------|------------------------------------------------|
| UserPrincipalName | unfamiliarsignin2@ntnuinft2504.onmicrosoft.com |
| CorrelationId | 9c306d0b-e030-404c-9f6c-9c4aeaa260b3 |
| AppDisplayName | Azure Portal |
| UserId | 7f33db1d-899c-4cae-b114-8108d96a7b3d |
| Type | SigningLogs |
| FailureCountBeforeSuccess | 37 |
| StartTime [UTC] | 2024-04-16T12:40:00Z |
| EndTime [UTC] | 2024-04-10T13:00:00Z |
| IPAddress | 62.148.32.132 |
| set_Browser | ["Safari 16.4"] |
| set_City | ["Fornebu"] |
| set_State | ["Akershus"] |
| set_Region | ["NO"] |
| set_ResultType | ["50074", "50097", "500121",] |
| UserPrincipalName1 | unfamiliarsignin2@ntnuinft2504.onmicrosoft.com |
| avgFailures | 0.9285714285714286 |
| Deviation | 38.84615384615384 |
| timestamp [UTC] | 2024-04-16T12:40:00Z |
| Name | unfamiliarsignin2 |
| UPNSuffix | ntnuinft2504.onmicrosoft.com |

Table 4.8: Analytics Rule: *Brute force attack against Azure Portal*

From table 4.8 we are given three results under the `set_ResultType`. These results indicate what went wrong during the attack, which triggered the analytics rule and created an incident in our system. Summarised, the errors `50074`, `50097` and `500121` mean that the system required authentication in the form of MFA,

but this was never given, or accepted, from the user [109][110].

Since the user had to write in a number to accept the MFA challenge this attack is considered to be failed, as described in 3.2.1.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

This attack was not detected in ID Protection.

Due to the user having to write in a response to the MFA prompt, this attack can be considered as not successful (3.2.1).

The attack can also be seen as mitigated, since the attacker did not gain access to the system, due to having to write in the right number from the MFA challenge. This mitigation came from the *per-user MFA*.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Yes*

Results with Conditional Access Enabled

For both attempts to sign in to `myapps.microsoft.com` and the *Azure portal*, these three CAPs were applied (see figure 4.10 for more details about the policies result status):

- *Require multifactor authentication for all users*
- *No persistent browser session*
- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users*

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details Conditional Access Report-only ...

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ | |
|------------------------------------------------------------------|-----------------------------|-------------------------------|-----------|-----|
| Require multifactor authentication for all users | Require multifactor auth... | | Failure | ... |
| No persistent browser session | | Sign-in frequency, Persist... | Success | ... |
| Require compliant or hybrid Azure AD joined... | Require multifactor auth... | | Failure | ... |

Figure 4.10: From ME-ID Sign-in logs: All the CAPs that were applied after denying the MFA challenge

The same CAPs were applied and resulted in success once the conditions of MFA were met (4.11).

The screenshot shows a table titled 'Activity Details: Sign-ins' with a search bar and several tabs: Basic info, Location, Device info, Authentication Details, Conditional Access (selected), and Report-only. The table lists three Conditional Access policies that were applied, all resulting in 'Success'.

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ |
|-----------------------------------|-----------------------------------|------------------------------------|-----------|
| Require multifactor authentica... | Require multifactor authentica... | | Success |
| No persistent browser session | | Sign-in frequency, Persistent b... | Success |
| Require compliant or hybrid A... | Require multifactor authentica... | | Success |

Figure 4.11: From ME-ID Sign-in logs: All the CAPs that were applied after completing the MFA challenge

Due to the success condition outlined in 3.2.1, this attack can not be considered a success. The attack was mitigated because of the MFA challenges being failed, which denied the attacker access.

CA being enabled did not trigger any incidents in Sentinel or any risk detections in ID Protection for *myapps*.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Yes*

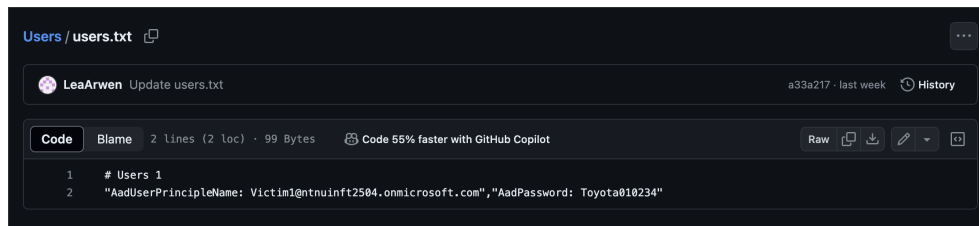
For *Azure Portal* the same incident was triggered in Sentinel, but no risk detections in ID Protection.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Yes*

4.4 Gather Victim Info

After waiting for more than 48 hours, no detections had been triggered and the .txt file containing user credentials (4.12) had not been discovered by Microsoft.



The screenshot shows a GitHub repository interface for a file named 'Users / users.txt'. The file was updated by 'LeaArwen' a39a217 last week. The file content is as follows:

```
1 # Users 1
2 "AadUserPrincipalName: Victim1@ntnuinf2504.onmicrosoft.com", "AadPassword: Toyota010234"
```

Figure 4.12: .txt file uploaded to a public GitHub repository containing user credentials (see 3.2.1).

Due to the nature of the test, it does not have a success condition as described in 3.2.1. The attack did not trigger any incidents in Sentinel, no risk detection in ID Protection, and no CAPs were applied. There has also been no sight of anyone attempting to sign in to the user in the ME-ID Sign-in logs.

This can be summarised:

1. Was the attack successful? *Not applicable*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

4.5 Proxy

4.5.1 Proxy (without MFA)

Results from Sentinel

This attack was not detected in Sentinel. However, the attack can be seen as successful because it fulfilled the success condition outlined in 3.2.1. This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

The attack was detected by ID Protection and a risk detection appeared in the ID Protection dashboard.

A user was registered as high-risk, and the detection type was registered as *Anonymous IP address*. Both sign-in and detection times were registered at 10:03, with the detection timing as real-time. The attack type was registered as *Obfuscation/Access using proxy, Access using a valid account (Detected Offline)* (see figure 4.13).

Based on the success condition in 3.2.1 this attack can be seen as successful and the attack was not mitigated as the attacker was granted access.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *No*

| Risk Detection Details | |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| User's risk report User's sign-ins User's risky sign-ins Linked risky sign-in User's risk detections | |
| Detection type | Anonymous IP address ⓘ |
| Risk state | Dismissed |
| Risk level | High |
| Risk detail | Admin dismissed all risk for user |
| Attack type(s) | Obfuscation/Access using proxy, Access using a valid account (Detected Offline) |
| Source | Identity Protection |
| Detection timing | Real-time |
| Activity | Sign-in |
| Detection time | 4/9/2024, 10:03 AM |
| Detection last updated | 4/9/2024, 10:15 AM |
| Token issuer type | Microsoft Entra ID |
| Sign-in time | 4/9/2024, 10:03 AM |
| IP address | 51.158.115.62 |
| Sign-in location | Paris, Paris, FR |
| Sign-in client | Mozilla/5.0 (Windows NT 10.0; rv:109.0) Gecko/20100101 Firefox/115.0 |
| Sign-in request id | e64f9b9d-6f4a-49f2-9faa-9beba0b86000 |
| Sign-in correlation id | 712dadaf-baa8-4143-aa8d-5d44ee3208e2 |

Figure 4.13: Risky User (without MFA) detected in ID Protection, after risk has been dismissed by admin.

Results with Conditional Access Enabled

The user was unable to sign in without MFA making the attack unsuccessful, based on the success condition outlined in 3.2.1. The user was registered not having MFA, as seen in activity details 4.15. Thereby access was denied to the user as seen in figure 4.16, and the attack was mitigated.

These five CAPs were applied. The conditions of MFA were not met, thereby failing the policies:

- *Require multifactor authentication for admins* - Status: Failure
- *Require multifactor authentication for all users* - Status: Failure
- *Require multifactor authentication for risky sign-ins* - Status: Failure
- *No persistent browser session* - Status: Success
- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users* - Status: Failure

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details **Conditional Access** Report-only ⋮

🔍 Search

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ |
|-----------------------------------------------------------------------|----------------------------------|-------------------------|----------------------------------------------|
| Require multifactor authentication for admins | Require multifactor authentic... | | Failure ⋮ |
| Require multifactor authentication for all users | Require multifactor authentic... | | Failure ⋮ |
| Require multifactor authentication for risky sign-ins | Require multifactor authentic... | | Failure ⋮ |
| No persistent browser session | | Sign-in frequency, P... | Success ⋮ |
| Require compliant or hybrid Azure AD joined devi... | Require multifactor authentic... | | Failure ⋮ |

Figure 4.14: From ME-ID Sign-in logs: All the CAPs that were applied and their result status

Activity Details: Sign-ins ×

Basic info Location Device info **Authentication Details** Conditional Access Report-only ⋮

Authentication Policies Applied **Session Lifetime Policies Applied**

Conditional Access Sign-in frequency (periodic re-authentication)
Identity Protection

| Date | Authentication met... | Authentication met... | Succeeded | Result detail |
|-----------------------|-----------------------|-----------------------|-----------|------------------------------------------------|
| 4/19/2024, 2:39:06 PM | Password | Password in the cloud | false | User needs to complete Multi-factor authenti.. |

Figure 4.15: From ME-ID Sign-in logs: Authentication method failing

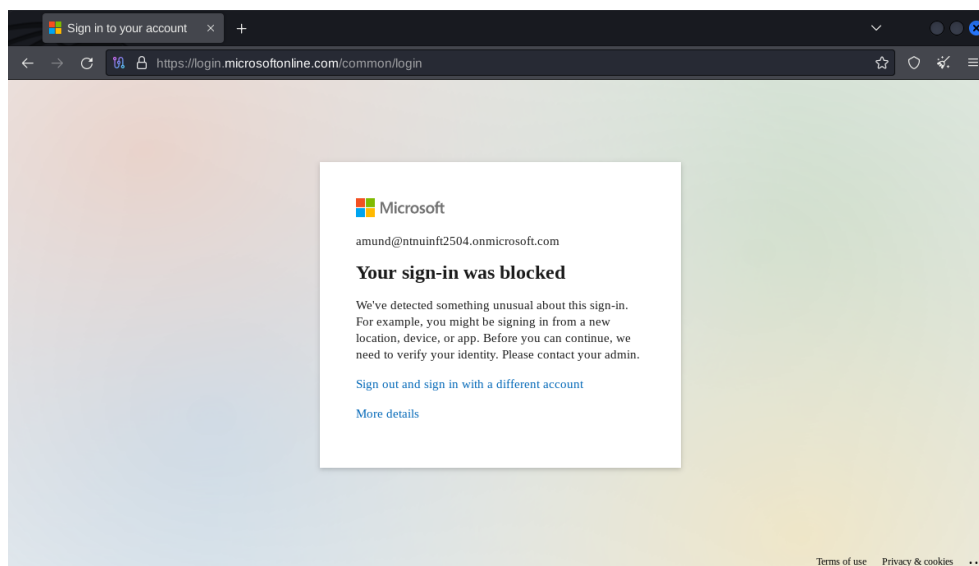


Figure 4.16: User attempting to sign in using proxy without MFA

CA being enabled did not trigger any new incidents in Sentinel and the same risk detection in ID Protection was registered as when CA was disabled.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Yes*

4.5.2 Proxy (with MFA)

Results from Sentinel

This attack was not detected in Sentinel. However, the attack can be seen as successful based on the success condition outlined in 3.2.1.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

The attack was detected in ID Protection. A risk detection was registered with the detection type *Anonymous IP address* as seen in risk details (figure 4.17). The user was registered as having passed MFA, which led to the user being remediated and thereby no notification of a risky user was issued. There were no further mitigating steps that prevented access, therefore we consider this attack to not have been mitigated.

Based on the success condition outlined in 3.2.1, this attack was successful. This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *No*

Risk Detection Details [Close]

[User's risk report](#)
[User's sign-ins](#)
[User's risky sign-ins](#)
[Linked risky sign-in](#)
[User's risk detections](#)

Detection type Anonymous IP address ⓘ
Risk state Remediated
Risk level Medium
Risk detail User passed multifactor authentication
Attack type(s) Obfuscation/Access using proxy, Access using a valid account (Detected Offline)
Source Identity Protection
Detection timing Real-time
Activity Sign-in
Detection time 4/10/2024, 1:11 PM
Detection last updated 4/10/2024, 1:13 PM
Token issuer type Microsoft Entra ID
Sign-in time 4/10/2024, 1:11 PM
IP address 195.160.220.104
Sign-in location Kyiv, Kyiv Misto, UA
Sign-in client Mozilla/5.0 (Windows NT 10.0; rv:109.0) Gecko/20100101 Firefox/115.0
Sign-in request id 3a200259-1f86-449d-a6b5-06c3dabd7900
Sign-in correlation id 3415323b-f099-450d-942a-0aad8c101132

Figure 4.17: Risky User detected in ID Protection (with MFA)

Results with Conditional Access Enabled

Due to the condition of MFA being met (see figure 4.18), the CAPs resulted in success. The attacker was given access to sign-in, thus the attack was not mitigated. These are the CAPs that were applied and their result status:

- *Require multifactor authentication for all users* - Status: Success
- *Require multifactor authentication for risky sign-ins* - Status: Success
- *No persistent browser session* - Status: Success
- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users* - Status: Success

Activity Details: Sign-ins [Close]

[Basic info](#)
[Location](#)
[Device info](#)
[Authentication Details](#)
[Conditional Access](#)
[Report-only](#)
⋮

Authentication Policies Applied **Session Lifetime Policies Applied**
 Conditional Access Sign-in frequency (periodic re-authentication)
 Identity Protection

| Date | Authentication met... | Authentication met... | Succeed... | Result detail | Requirement |
|-----------------------|-------------------------|-----------------------|------------|---------------------------|-------------|
| 4/19/2024, 2:24:20 PM | Password | Password in the cloud | true | Correct password | |
| 4/19/2024, 2:24:20 PM | Mobile app notification | | true | MFA completed in Azure AD | |

Figure 4.18: From ME-ID Sign-in logs: Authentication method being succeeded

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details **Conditional Access** Report-only ...

🔍 Search

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ |
|-----------------------------------------------------------------------|------------------------------------|---------------------------|-----------|
| Require multifactor authentication for all users | Require multifactor authentication | | Success |
| Require multifactor authentication for risky sign-ins | Require multifactor authentication | | Success |
| No persistent browser session | | Sign-in frequency, Per... | Success |
| Require compliant or hybrid Azure AD joined devi... | Require multifactor authentication | | Success |

Figure 4.19: From ME-ID Sign-in logs: All the CAPs that were applied and their result status

CA being enabled did not trigger any new incidents in Sentinel and the same risk detection in ID Protection was registered as when CA was disabled.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *No*

4.6 Brute Force

4.6.1 Brute Force: Password Guessing

Results from Sentinel

This attack was not detected in Sentinel. This attack can be seen as unsuccessful, as it did not fulfil the success condition described in 3.2.1.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *No*
3. Was the attack mitigation? *Not applicable*

Results from ID Protection

This attack was not detected in ID Protection.

Due to the success condition outlined in section 3.2.1 not being met, this attack was unsuccessful.

This attack was mitigated due to the attacker being prompted with MFA, and also due to the user account being locked by the *Azure smart lockout*.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *No*

3. Was the attack mitigated? *Yes*

Results with Conditional Access enabled

This attack was not detected while CA was enabled. No incidents were triggered in Sentinel and no risk detections in ID Protection.

The attacker was not able to guess the right credentials before the account got locked and connection-rate locked by *Azure smart lockout* where the default amount of attempts before lockout is set to 10 (2.2.4). When the correct credentials were later attempted, we received the same message as before, due to *Azure smart lockout*. Considering the success condition from 3.2.1, this would not entice a success and the attack is considered mitigated.

This can be summarised:

1. Was the attack successful? *No*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Yes*

4.6.2 Brute Force: Password Spray

Results from Sentinel

When we carried out this attack, described in 3.2.1, Sentinel created *one* incident from the configured analytics rule *Password spray attack against Microsoft Entra ID application*. The result for the incident is described in table 4.9 provided below.

| | |
|-----------------------------|------------------------------------------|
| IPAddress | 92.127.56.129 |
| StartTime [UTC] | 2024-04-15T14:01:31.8105416Z |
| EndTime [UTC] | 2024-04-16T13:30:33.3712243Z |
| TargetedApplication | Azure Active Directory PowerShell |
| FailedPrincipalCount | 11 |
| UserPrincipalNames | restuser10@ntnuinft2504.onmicrosoft.com |
| | restuser11@ntnuinft2504.onmicrosoft.com |
| | restuser8@ntnuinft2504.onmicrosoft.com |
| | leakeduser1@ntnuinft2504.onmicrosoft.com |
| | restuser6@ntnuinft2504.onmicrosoft.com |
| | leakeduser2@ntnuinft2504.onmicrosoft.com |
| | restuser9@ntnuinft2504.onmicrosoft.com |
| | restuser2@ntnuinft2504.onmicrosoft.com |
| | restuser7@ntnuinft2504.onmicrosoft.com |
| | restuser1@ntnuinft2504.onmicrosoft.com |
| | admin1@ntnuinft2504.onmicrosoft.com |
| UserDisplayNames | Rest User 10 |
| | Rest User 11 |
| | Rest User 8 |
| | Leaked User 1 |
| | Rest User 6 |
| | Leaked User 2 |
| | Rest User 9 |
| | Rest User 2 |
| | Rest User 7 |
| | Rest User 1 |
| | Admin 1 |
| ClientAppsUsed | Mobile Apps and Desktop Clients |
| Locations | NO Oslo Oslo |
| FailureCountByPrincipal | [21,20,9,9,9,9,9,9,9,9] |
| WindowsThresholdBreaches | 3 |
| Type | SignInLogs |
| Type1 | SignInLogs |
| GlobalSuccessPrincipalCount | 2 |
| ResultTypeSuccesses | ["50076"] |
| GlobalFailPrincipalCount | 11 |
| ResultTypeFailures | ["50126","50053"] |
| timestamp [UTC] | 2024-04.15T14:02:31.8105416Z |

Table 4.9: Analytics Rule: Password spray attack against Microsoft Entra ID application

From table 4.9 we are given *three* results, *one* in the *ResultTypeSuccesses* and

two in the *ResultTypeFailures*.

The result for success, *50076 - UserStrongAuthClientAuthNRequired*, which means that MFA was required during this sign-in [110], shows that the attack succeeded because the right credentials were used for a user, meeting the success condition mentioned in section 3.2.1.

The two failure results indicate what errors were registered during the attack. *50126 - InvalidUserNameOrPassword* and *50053* indicate that the wrong password was used and that the user was locked as a response to the many sign-in attempts [110].

Since we got informed of needing MFA the result can be seen as successful, as this indicated that we had used the right credentials for a user.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

This attack was not detected in ID Protection, nor mitigated. Due to meeting the success condition outlined in section 3.2.1 this attack can be considered successful.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

Results with Conditional Access Enabled

Due to the condition of MFA not being met, the CAPs failed. The attack was mitigated as the attacker was denied access. These are the three CAPs that were applied and their result status (see figure 4.20):

- *Require multifactor authentication for all users* - Status: Failure
- *No persistent browser session* - Status: Success
- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users* - Status: Failure

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details **Conditional Access** Report-only ...

🔍 Search

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ | |
|------------------------------------------------------------------|---------------------------------------|-------------------------------|-----------|-----|
| Require multifactor authentication for all users | Require multifactor authentication... | | Failure | ... |
| No persistent browser session | | Sign-in frequency, Persist... | Success | ... |
| Require compliant or hybrid Azure AD joined... | Require multifactor authentication... | | Failure | ... |

Figure 4.20: From ME-ID Sign-in logs: All the CAPs that were applied

The attacker was able to gain access to a user's credentials, but the sign-in attempt was not successful due to CA's requirement of MFA. Based on the attack's success condition in 3.2.1, this can be considered a success.

The attack was still detected with Sentinel, CA being enabled did not trigger any new incidents in Sentinel or risk detections in ID Protection.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Yes*

4.7 Steal Web Session Cookie

4.7.1 Steal Web Session Cookie (same IP address)

Results from Sentinel

This attack was not detected in Sentinel. The attack is considered successful, due to the fulfilment of the success condition outlined in 3.2.1.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

This attack was not detected in ID Protection, nor mitigated. Due to meeting the success conditions outlined in 3.2.1, this attack can be considered successful.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

Results with Conditional Access Enabled

The CAPs resulted in success due to the condition of MFA being met, the authentication details also show that the authentication method has previously been satisfied, as seen in figure 4.21. The attacker was then able to sign in, thereby the attack was not mitigated.

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details Conditional Access Report-only ...

| Authentication Policies Applied | | Session Lifetime Policies Applied | | | |
|---------------------------------|-----------------------|------------------------------------------------|-----------|----------------------------|-------------|
| Conditional Access | | Sign-in frequency (periodic re-authentication) | | | |
| Date | Authentication met... | Authentication met... | Succeeded | Result detail | Requirem... |
| 4/19/2024, 1:52:31 PM | Previously satisfied | | true | First factor requiremen... | |
| 4/19/2024, 1:52:31 PM | Previously satisfied | | true | MFA requirement satis... | |

Figure 4.21: From ME-ID Sign-in logs: Authentication method being previously satisfied

These are the three policies that were applied, (see figure 4.22):

- *Require multifactor authentication for all users* - Status: Success
- *No persistent browser session* - Status: Success
- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users* - Status: Success

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details Conditional Access Report-only ...

🔍 Search

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ | |
|------------------------------------------------------------------|------------------------------|-------------------------------|-----------|-----|
| Require multifactor authentication for all users | Require multifactor authe... | | Success | ... |
| No persistent browser session | | Sign-in frequency, Persist... | Success | ... |
| Require compliant or hybrid Azure AD joined... | Require multifactor authe... | | Success | ... |

Figure 4.22: From ME-ID Sign-in logs: All the CAPs that were applied and their result status

CA being enabled did not help trigger any new incidents in Sentinel or any risk detections in ID Protection.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

4.7.2 Steal Web Session Cookie (new IP address)

Results from Sentinel

This attack was not detected in Sentinel. The attack is considered successful, due to the fulfilment of the success condition outlined in 3.2.1.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

This attack was not detected in ID Protection, nor mitigated. Due to meeting the success condition outlined in 3.2.1, this attack can be considered successful.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

Results with Conditional Access Enabled

The CAPs resulted in success due to the condition of MFA being met, thereby giving the attacker access to sign in using the web session cookie. The authentication details also show that the authentication method has previously been satisfied (see figure 4.23). The attack was not mitigated.

| Activity Details: Sign-ins | | | | | |
|----------------------------------------|-----------------------|------------------------------------------------|------------------------|----------------------------|-------------|
| Basic info | Location | Device info | Authentication Details | Conditional Access | Report-only |
| Authentication Policies Applied | | Session Lifetime Policies Applied | | | |
| Conditional Access | | Sign-in frequency (periodic re-authentication) | | | |
| Date | Authentication met... | Authentication met... | Succeeded | Result detail | Requirem... |
| 4/19/2024, 1:52:31 PM | Previously satisfied | | true | First factor requiremen... | |
| 4/19/2024, 1:52:31 PM | Previously satisfied | | true | MFA requirement satis... | |

Figure 4.23: From ME-ID Sign-in logs: Authentication method being previously satisfied

These are the three policies applied and can also be seen in figure 4.24:

- *Require multifactor authentication for all users* - Status: Success
- *No persistent browser session* - Status: Success
- *Require compliant or hybrid Azure AD joined device or multifactor authentication for all users* - Status: Success

Activity Details: Sign-ins ×

Basic info Location Device info Authentication Details **Conditional Access** Report-only ...

Search

| Policy Name ↑↓ | Grant Controls ↑↓ | Session Controls ↑↓ | Result ↑↓ | |
|------------------------------------------------------------------|---------------------------------------|-------------------------------|-----------|-----|
| Require multifactor authentication for all users | Require multifactor authentication... | | Success | ... |
| No persistent browser session | | Sign-in frequency, Persist... | Success | ... |
| Require compliant or hybrid Azure AD joined... | Require multifactor authentication... | | Success | ... |

Figure 4.24: From ME-ID Sign-in logs: All the CAPs that were applied and their result status

CA being enabled did not trigger any incidents in Sentinel, but it did trigger a new risk detection in ID Protection.

ID Protection after turning on Conditional Access

The attack was detected in ID Protection once CAPs were enabled.

The user was registered signing in at 1:52 PM, and the detection time was at 9:24 PM. The detection type *Anomalous token* was detected, with the attack types being *Access using a valid account (Detected Offline)* and *Steal Web Session Cookie/Token Theft*, as seen in *attack type(s)* (figure 4.25).

Risk Detection Details ×

👤 User's risk report
 🔄 User's sign-ins
 🚨 User's risky sign-ins
 🔗 Linked risky sign-in
 🚩 User's risk detections

| | |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| Detection type | Anomalous token ⓘ |
| Risk state | At risk |
| Risk level | Medium |
| Risk detail | - |
| Attack type(s) | Access using a valid account (Detected Offline), Steal Web Session Cookie/Token Theft |
| Source | Identity Protection |
| Detection timing | Offline |
| Activity | User |
| Detection time | 4/19/2024, 9:24 PM |
| Detection last updated | 4/19/2024, 9:24 PM |
| Token issuer type | Microsoft Entra ID |
| Sign-in time | 4/19/2024, 1:52 PM |
| IP address | 45.130.203.77 |
| Sign-in location | Al Qahirah, Al Qahirah, EG |
| Sign-in client | Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/124.0.0.0 Safari/537.36 Edg/124.0.0.0 OS/10.0.22631 |
| Sign-in request id | b1530f73-955e-4fdc-88ee-f80106b51a00 |

Figure 4.25: From ID Protection Risk Detection Details: Details of the risk detected

The attacker was granted access through the reuse of the stolen web session cookie, making this attack a success based on the success condition presented in 3.2.1.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *No*

4.8 Account Manipulation

4.8.1 Add Cloud Roles

Results from Sentinel

When we carried out this attack, described in 3.2.1, Sentinel created *two* incidents triggered by the analytics rules *New User Assigned to Privileged Role* and *User Assigned New Privileged Role*. The result for each incident is described in table 4.10 and table 4.11 provided down below.

| | |
|---------------------------------|----------------------------------------|
| TimeGenerated [UTC] | 2024-04-09T08:39:09.6661973Z |
| OperationName | Add eligible member to role |
| RoleName | Global Administrator |
| Target | RestUser5@ntnuinft2504.onmicrosoft.com |
| Initiator | MS-PIM |
| InitiatingAppName | MS-PIM |
| InitiatingAppServicePrinicpalId | 8981fda0-598e-4722-8a93-346a5a228977 |
| Result | Success |
| TargetName | RestUser5 |
| TargetUPNSuffix | ntnuinft2504.onmicrosoft.com |

Table 4.10: Analytics Rule: *New User Assigned to Privileged Role*

| | |
|---------------------------------|----------------------------------------------------------|
| TimeGenerated [UTC] | 2024-04-09T08:00:00Z |
| OperationName | Add eligible member to role in PIM requested (permanent) |
| RoleName | Global Administrator |
| Target | RestUser5@ntnuinft2504.onmicrosoft.com |
| Initiator | lea@ntnuinft2504.onmicrosoft.com |
| InitiatingAppName | lea@ntnuinft2504.onmicrosoft.com |
| InitiatingAppServicePrincipalId | 90f53014-b7e3-4dfe-b91c-6927384f673c |
| Result | Success |
| TargetName | RestUser5 |
| TargetUPNSuffix | ntnuinft2504.onmicrosoft.com |
| InitiatorName | lea |
| InitiatorUPNSuffix | ntnuinft2504.onmicrosoft.com |

Table 4.11: Analytics Rule: *User Assigned New Privileged Role*

Due to the fulfilment of the success condition, outlined in 3.2.1, this attack can be seen as successful.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *Yes*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

When we carried out this attack, outlined in 3.2.1, ID Protection did not detect the attack.

Due to meeting the success condition outlined in section 3.2.1 this attack can be considered successful.

This attack was not mitigated by ID Protection.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

Results with Conditional Access Enabled

Due to the nature of this test, no CAPs were applied. However, the attack was still detected in Sentinel with the same incidents being triggered, while no risk detections were triggered in ID Protection. This attack can be considered successful as the success condition outlined in 3.2.1 were met. The attack was not mitigated.

This can be summarised:

1. Was the attack successful? *Yes*

2. Was the attack detected? *Yes*
3. Was the attack mitigated? *No*

4.8.2 Add Cloud Credentials

Results from Sentinel

This attack was not detected in Sentinel. The attack can, however, be considered successful, due to the fulfilment of the success condition outlined in 3.2.1.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

This attack was not detected nor mitigated in ID Protection. Due to meeting the success condition outlined in section 3.2.1 this attack can be considered successful.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

Results with Conditional Access Enabled

This attack was not detected while CA was enabled, nor mitigated. No incidents in Sentinel were triggered or any risk detections in ID Protection. Due to meeting the success condition outlined in section 3.2.1 this attack can be considered successful.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

4.9 Create Account

Results from Sentinel

This attack was not detected in Sentinel. However, since the attacker fulfilled the success condition of making an account, the attack can be seen as successful, as outlined in 3.2.1.

This can be summarised:

1. Was the attack successful? *Yes*

2. Was the attack detected? *No*
3. Was the attack mitigated? *Not applicable*

Results from ID Protection

This attack was not detected in ID Protection, nor mitigated.

Due to meeting the success condition outlined in section 3.2.1 this attack can be considered successful.

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

Results with Conditional Access Enabled

This attack was not detected while CA was enabled, nor mitigated. Due to meeting the success condition outlined in section 3.2.1 this attack can be considered successful. It did not trigger any incidents in Sentinel or any risk detections in ID Protection

This can be summarised:

1. Was the attack successful? *Yes*
2. Was the attack detected? *No*
3. Was the attack mitigated? *No*

Chapter 5

Discussion

5.1 Outline of Chapter

In this chapter, the results will be discussed according to the sub-research questions presented in 1.3. First, we will discuss the answers to each of the sub-questions before answering the main research question. This answer will be based on the conclusions to the sub-questions.

5.2 Sub-Research Question 1

In this part, we will consider the first sub-research question (1.3):

1. *How does Sentinel, configured with the rulesets for Entra ID and Entra ID Protection provided by Microsoft in Content Hub, provide any additional security features which are not available through the Entra ID Protection dashboard?*

To answer this question, we will utilise the theory regarding the Sentinel dashboard, as presented in section 2.5.6, and the theory regarding the ID Protection dashboard, as presented in 2.3.3, to compare the two dashboards. Additionally, we will examine our test results to discern if there are any discrepancies between what attacks were detected by each system. With this method, we will be able to see what security measures were provided in the different dashboards.

The Overview Dashboards

The overview dashboard provided in both Sentinel, see figure 5.1, and ID Protection, see figure 5.2, are quite similar. They both provide you with some information about an incident or an attack¹ that was detected. Since these dashboards are meant to provide an overview, they only offer general information. You can investigate further by accessing the widgets presented in the overview dashboards.

¹Incident in Sentinel, attack in ID Protection.

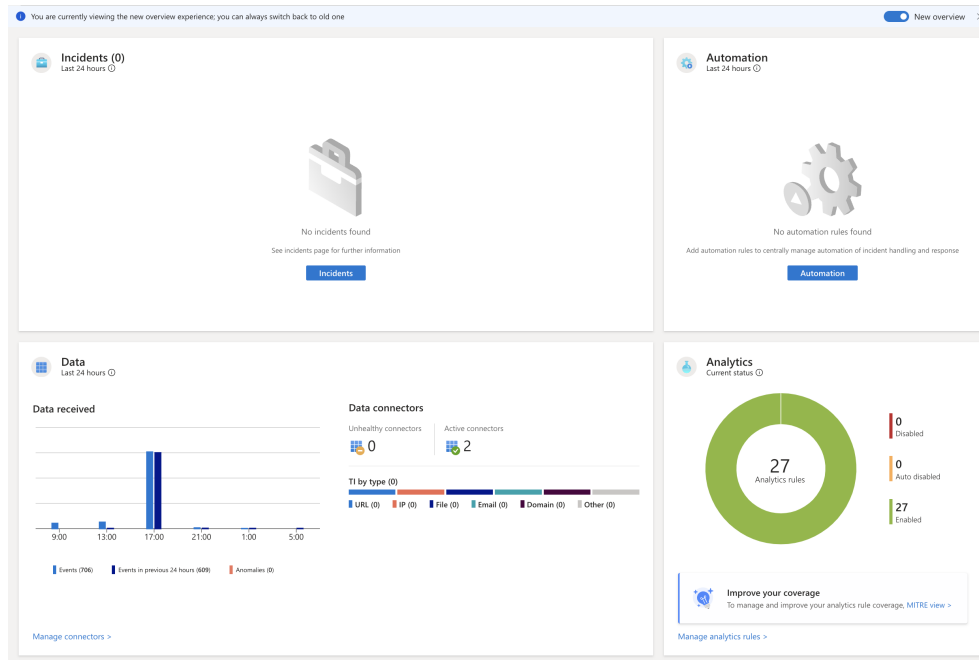


Figure 5.1: Sentinel overview dashboard

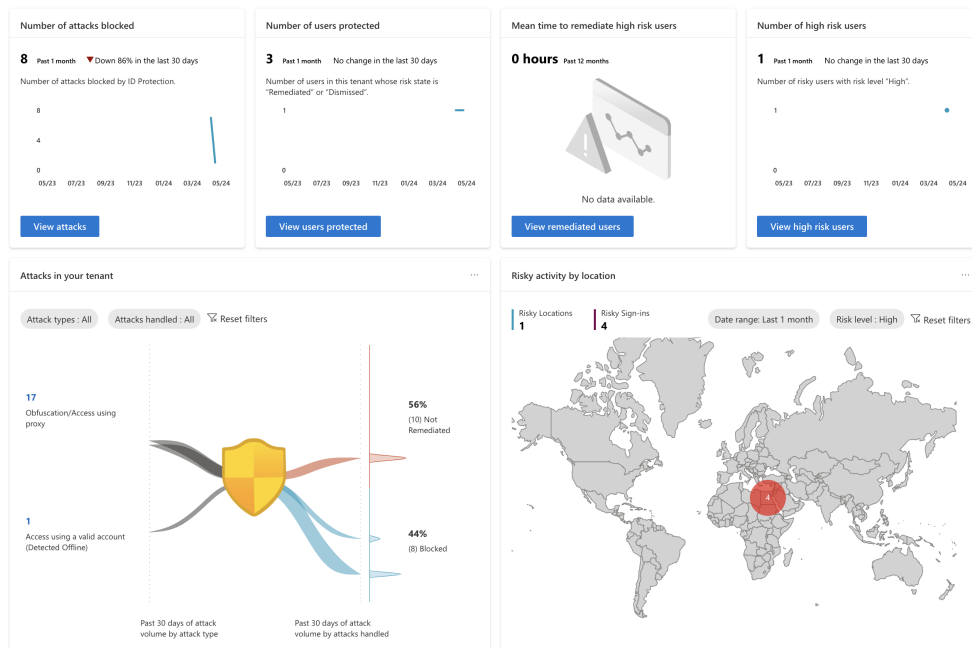


Figure 5.2: ID Protection overview dashboard

However, there are some differences between the overview dashboards and the information they provide. In the ID Protection dashboard you receive inform-

ation about *how many* in various categories, such as how many blocked attacks, how many users have been protected, how many attacks in your tenant and such. In some cases, you are also provided with a graph which indicates the occurrences over the past month. However, you do *not* receive any specific details about the attacks, protected users, or blocked attacks. Practically, this means that you are presented with a number indicating how many times something has happened, but you are not presented with any other information about what has happened. To get more information than just a number, you need to enter the different widgets dashboard specifically, where you can receive better information about a specific attack.

In the Sentinel overview dashboard, you are similarly not provided with detailed event information. However, here you receive more categorised information, see section 2.5.6. This additional categorisation may offer clearer insights, compared to the ID Protection dashboard, helping you understand if proactive actions are necessary, such as addressing an unhealthy data connector, or if reactive measures are required for incidents with *High* severity.

Incidents Dashboard vs. Risk Detections Dashboard

In our thesis, we are focusing on the detection of attacks regarding user identities. Therefore, it is natural for us to compare the incidents dashboard from Sentinel, see figure 5.3, to the Risk detections dashboard, see figure 5.4, in ID Protection. Both of these dashboards are widgets in the overview dashboard, and one can find more information about the detected attacks by clicking on them.

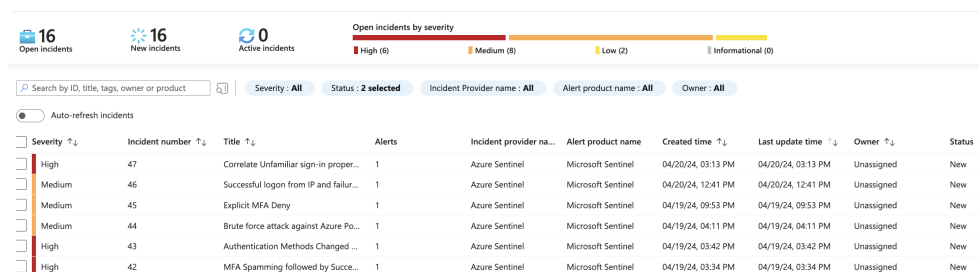


Figure 5.3: Sentinel incident dashboard




Figure 5.4: ID Protection risk detections dashboard

These dashboards are similar, both present a list of incidents (detected attacks) within a given time window. The most noticeable difference is that in the Sentinel

Incidents dashboard, you are given a summary of how many incidents you have open, how many new detections there are, and how many of these are active. Additionally, a graph is provided to illustrate the severity of the open incidents and the distribution of incidents across severity levels. In the Risk Detection dashboard in ID Protection you are not given any summary of the detected attacks, you are only provided with the list of detected Risk Detections.

The list of detected incidents and risk detections provides similar information about each event. One of the differences between the given information is what the different columns of categories are called. For example, while the incident dashboard uses the term *severity* the risk detection dashboard uses the term *risk level*. In this example, it is also noticeable that in Sentinel, the severity is also indicated with colours; red for *High*, orange for *Medium* and yellow for *Low*, which makes it easy to get a fast overview of the severity of the incidents in the list.

Another thing to notice about the Incident and Risk Detections dashboards is that the Risk Detections dashboard only provides you with information specific to the event that triggered the detection, such as detection time, user, location, detection type and more. In the Incident dashboard, you are also given information about the triggering event, such as severity, title, and creation time. However, it appears that this dashboard is designed more to provide an overview of the detected incidents, and how they will be handled, as you receive information about the status, owner, updated time, and tags.

Findings from our Results

When simulating attacks and collecting our results, presented in section 4, we gained an overview of the differences between what the incident dashboard in Sentinel presented and what the risk detection dashboard in ID Protection presented. As these dashboards are based on different detection methods - analytics rules and risk detections - we were expecting a difference in the number of attack techniques each system was able to detect, but the difference was greater than expected.

The first notable observation would be the number of attacks detected by each system. Of the 15 attacks we performed, six of them were detected by Sentinel, which resulted in the creation of nine incidents, with some attacks triggering multiple rules. Only three of the attacks were detected by ID Protection². Of these attacks, only one attack, *Valid account: New location*, was detected by both Sentinel and ID Protection. This means that Sentinel detected five attacks that ID Protection did not detect, and ID Protection discovered two attacks which Sentinel did not detect. Here we expected more attacks to be detected by both security tools.

The two following lists will state which attacks Sentinel discovered and ID Protection did not discover, and vice versa.

²These numbers are collected from section 4, and are equal to where the answer *Yes* is given to the question *Was the attack detected?*

Only Sentinel:

- Valid Account: New device
- MFA request (explicit deny)
- MFA request (no answer)
- Brute Force: Password Spray
- Add cloud roles

Only ID Protection:

- Proxy (without MFA)
- Proxy (with MFA)

The difference in the number of detected attacks shows that Sentinel provided some additional security during our testing, that was not provided through the ID Protection dashboard. However, as mentioned previously, we expected a difference, but we were surprised by how large the difference was.

Both Sentinel and ID Protection have detection methods for detecting the password spray attack technique (see section 2.3.2 and table 3.1). Therefore, we were expecting this attack to be detected by both tools. However, from our results, we see that only Sentinel detected the attack. Furthermore, we expected the test *Valid Account: New device* to also trigger the *Unfamiliar sign-in properties* risk detection in ID Protection due to the test following Microsoft's guide for simulating risk detections (see 3.2.1).

The reason for this difference might be that ID Protection uses a machine-learning algorithm to detect the risk detections. This method helps reduce false positives, as mentioned in section 2.3.4. Since we performed most of our attacks on the same devices and the same IPs (see 3.6.1), it is possible that ID Protection did label the events as false positives which could explain why we were provided with so few risk detections.

Since the incidents in Sentinel are based on the enabled analytics rules from Content Hub, these incidents will be created as long as the conditions in the analytics rules are met. This means that Sentinel will not remove any false positives if this is not defined within the rule itself (see 2.5.4), as this depends on how the configuration of the different analytics rules is handled.

Conclusion

There are quite a few differences between the dashboard in Sentinel and in ID Protection. These mechanisms do not detect the same things; one detects incidents based on a set of rules, whereas the other detects risk detections based on known attacks and a machine learning algorithm. With this sub-research question, we attempted to determine whether having Sentinel would provide any added value to an organisation's ability to detect security incidents, or if ID Protection would be enough.

What we have found is that the dashboards are similar but have some differences. While the dashboard in ID Protection provides more information about the risk detection itself, the incident dashboard in Sentinel gives you more alternatives to handle the incidents, making it a more proactive dashboard.

Another observation was that more attacks were detected and presented in Sentinel compared to in ID Protection. It is possible that this happened because of the machine learning behind the system, but it does still imply that ID Protection and Sentinel add value to one another.

In conclusion, we can confidently assert that Sentinel provides some additional security features not found in ID Protection. Utilising both systems will give an organisation more information about potential attacks, thereby allowing for a more secure detection system, compared to only having one of the systems.

5.3 Sub-Research Question 2

In this part, we will consider the second sub-research question (1.3):

2. *How can the rulesets in Sentinel be modified to reduce the number of false positive and false negative results?*

This sub-research question has two parts where the first is what changes to the rules are needed to lower the amount of false positive results and the second part considers how to lower the amount of false negative results. As mentioned in 3.1.1, we will find the answer to this sub-research question by evaluating a combination of the test results. These are described in 3.2. We will first discuss the false positives and then the false negative results.

5.3.1 Reducing False Positives

To reduce the number of false positives we will first define which of our results indicates a false positive result.

Based on the description of a false positive (see section 2.6), we define it as an analytics rule being triggered and having created an incident from an unsuccessful attack. Furthermore, it is important to note that in some cases of an unsuccessful attack, it is still necessary for the organisation to follow up on the event. An example of where an unsuccessful attack still could be nice to follow up on is where the attacker has the username and password of the user. In these cases, even if the attack is unsuccessful the attacker would have the credentials for the user and would be able to perform further attacks.

We can consider the following answers to our test questions to be the definition of a false positive:

1. Was the attack successful? **No**
2. Was the attack detected? **Yes**

Looking through our results, the following analytics rules triggered in cases with those test results:

Brute force attack against Azure Portal Test: MFA request against `portal.azure.com` without CAPs (4.3) - The attempted attack was detected, but unsuccessful. The correct password was used during sign-in, but the attack failed due to MFA being requested.

MFA explicit deny Test: MFA request (explicit deny) without CAPs (4.3) - The attempted attack was detected, but unsuccessful. The correct password was used during sign-in, but the attack failed due to MFA being requested.

However, both of these alerts require the attacker to have the user's password. Therefore further attacks can be performed, and even the same again, which means that they are not false positives. Therefore, we can conclude from our results that there were no false positive results. This warrants no changes to the analytics rules.

5.3.2 Reducing False Negatives

A false negative is the opposite of a false positive. In these cases, there was no detection, but still, the attack was successful (2.6). As with false positives, we are also counting unsuccessful attacks requiring an account's username and password as events that would need the organisation's handling.

This means that the following test results would indicate a false negative:

1. Was the attack successful? **Yes**
2. Was the attack detected? **No**

Of the attacks we simulated, multiple were successful. However, many did not create an incident in the Sentinel dashboard. From the results chapter we find that the following test had a false negative result:

- Valid Account: Atypical travel
- Gather Victim Info
- Proxy (without MFA)
- Proxy (with MFA)
- Steal Web Session Cookie (same IP)
- Steal Web Session Cookie (new IP)
- Add Cloud Credentials
- Create Account

Normally, using MFA would remediate the risk considered with a user and a given sign-in, proving that the attack was not dangerous (2.3.2). However, an attacker could utilise the *Steal Web Session Cookie* attack technique which allows an attacker to bypass the MFA challenge (2.7.3). Based on the results, this attack technique was not discovered (4.7). This could lead to an attacker being able to

perform further attacks, such as *Add Cloud Credentials* or *Create Account* - which also did not lead to any new incidents in Sentinel.

We will discuss these events more closely, starting with the attack techniques that have the goal of persistence within the cloud environment before assessing the attacks related to initial access. Lastly, we will discuss credential access. This investigation seeks to establish whether any of these cases warrants a change in the analytics rules, or if the present rules are sufficient.

Persistence

Starting with the attack technique *Add Cloud Credentials*, we see that none of these received any detection in Sentinel. There is an already existing analytics rule which attempts to detect this called *Suspicious sign-in followed by MFA modification* (A.2). The purpose of this rule is to create an incident if a user has had a suspicious sign-in followed by adding new MFA credentials. During our testing, we did not receive any new incidents, however, we used MFA when signing in. Because this rule uses UEBA, which is machine learning, we are back to the problems associated with machine learning. Microsoft gives us limited insight into the logic in use, making it difficult for us to perform a deep analysis of why the conditions of the rule were not satisfied (3.6.2). Therefore, we have to conclude that the rule probably does not need modification as long as it performs as advertised. However, it is important to ensure that any possibly suspicious sign-in events are still detected and reported.

There exists an analytics rule which detects the creation of new accounts, which was the goal of the attack technique *Create Account*. This rule is *Account created or deleted by non-approved user*. This rule checks whether the creator of an account is in a given list (A.2). The user we performed the test with was in this list, but it would be possible for the user to have been in use by a malicious actor. An organisation would be able to determine whether the user was used by a malicious actor by checking whether there was anything suspicious with their sign-in process. However, there are analytics rules which detect whether or not the user has had a suspicious sign-in event. The result of this is that if an incident was created because of a suspicious sign-in event for the user, the organisation would have been alerted of the security breach.

We can conclude that there is not needed any changes, or additions, to the analytics rule *account created or deleted by non-approved user*, which covers the attack technique *Create Account*. The reason for this is that there should be an incident created if the acting user is not in the list of approved users and if a compromised, but approved, user attempts to create a user, there should have been a detection earlier in the attack sequence. This brings us to the need for proper detection of attacks during their initial access step.

Initial Access

As concluded above, there is a need to discuss the false negative results related to the initial access attack vector, which in our case is found through sign-in activity. The attack techniques which yielded a false negative result were *Atypical travel* and *Proxy*.

Starting with the test for *Atypical travel* we see that there was no alarm created in Sentinel. The closest existing analytics rule which looks for atypical travel would be *Correlate unfamiliar sign-in properties and atypical travel alerts* (A.3). This rule creates an incident if there is detected both an *Unfamiliar sign-in properties* and a *Atypical travel* risk detection in the ID Protection dashboard during a short time frame of each other. This means that if there is not registered an unfamiliar sign-in properties event on the second location, no alert will be created in Sentinel.

A possible change would be to remove the need for the *Unfamiliar sign-in properties* risk detection. However, the benefit of needing the extra risk detection is that it would lower the chance of a false positive. An example of such a case would be the use of a VPN service, which would be a false positive result as long as there is no further suspicious behaviour. Therefore, only needing the *Atypical travel* risk detection could increase the amount of false positives.

There is another analytics rule which might be created in the case of *Atypical travel*. The *Successful logon from IP and failure from a different IP* rule detects whether there is a new sign-in from another IP and a second failed sign-in to the same Azure application from the same User within 10 minutes of each other (A.2). This rule did create an incident while testing with CAPs enabled, but not without which was unexpected as both tests were performed with MFA enabled.

There are no analytics rules which fit towards the *Proxy* attack technique. One method of fixing this hole in the security detection would be to create an analytics rule which would trigger if there was an *Anonymous IP address* risk detection triggered in ID Protection. This would create a direct link between what shows up in Sentinel and ID Protection, lowering the need for a security analyst to be working with two different tools at once to gain the appropriate coverage.

Credential Access

The most surprising result we found was that there were no incidents created in the Sentinel dashboard from the attack technique *Steal Web Session Cookie*. This attack is potentially very threatening as it enables the attacker to bypass the MFA challenge (2.7.3). During testing we found that there was no response to the reuse of a valid token on a new machine and a new IP address, however, there was a detection from ID Protection when CA was enabled (4.7.2). An appropriate analytics rule would be to detect the *Anomalous token* risk detection in ID Protection and create an incident in Sentinel based on risk detections of this type. This would decrease the need for a security specialist to address security incidents on both platforms simultaneously.

Additionally, because it is possible to bypass MFA by reusing a session cookie,

it could also be beneficial to check how the MFA challenge was met if the sign-in event happened from an unfamiliar location or unfamiliar device. If an attacker has stolen an active session cookie and is reusing it, the MFA challenge is labelled as being satisfied by *previously satisfied*. However, an analytics rule that creates an incident each time a sign-in happens with the MFA challenge satisfied by a *previously satisfied* would likely create many false positive events. This is because the SSO design of Azure would reuse the cookie when signing in to another application (2.1.2). If each of these events created a new incident the Sentinel dashboard would be crowded by unnecessary events. It is possible to modify the analytics rule to not create an incident for sign-in events which had sign-in properties that an organisation did not consider to be suspicious. However, the reuse of a token coupled with *Unfamiliar sign-in properties* is what prompts the creation of an *Anomalous token* risk detection (2.3.2), and as seen in our results, the attempts to utilise this attack technique were detected while CA was enabled. Therefore, we conclude that the evaluation of how the MFA challenge was met is not a property which alone warrants a change in the analytics rules. Still, we do see that a security analyst would need to evaluate whether this attack technique was utilised or not when assessing a suspicious sign-in incident, even if the MFA challenge was satisfied.

The last attack technique which did not create an incident was *Gather victim info*. This result is surprising as Microsoft has proven that a detection like this is possible. In ID Protection, there is both a risk detection which detects leaked credentials for user identities and workload identities (2.3.2). Microsoft recommends uploading workload credentials to GitHub to simulate the leaked credentials risk detection for workloads (2.3.4). Because of this recommendation, we deemed it likely that they also would have discovered user credentials which were uploaded to GitHub. However, this was not the case. These risk detections might have different searching methods which could yield different results, but due to these methods not being publicly available, we are not able to determine the exact reason.

The only detection which could be made in ME-ID and ID Protection that would indicate leaked user credentials would be the risk detection for *Leaked credentials* for users. Therefore, our conclusion is again to create a stronger bond between ID Protection and Sentinel and create an analytics rule which creates an incident for when this risk detection is triggered.

5.3.3 Conclusion

During testing, we found no false positives, but multiple false negatives. In general, the rulesets provided by Microsoft, through Content Hub, did detect most attacks which would be good to be alerted about. The changes we recommend are all additions intended to cover the attack techniques which there were no analytics rules focused on.

Here are the changes we recommend which were discussed above:

1. Create an analytics rule which creates an incident when the *anonymous IP address* risk detection has triggered in ID Protection.
2. Create an analytics rule which creates an incident when the *anomalous token* risk detection has triggered in ID Protection.
3. Create an analytics rule which creates an incident when the *leaked credentials* risk detection has triggered in ID Protection.

Of course, one would be able to use both the ID Protection dashboard and the Sentinel dashboard, but that would require the security analysts to keep up with two different services, and potentially increase the chance for alert fatigue (2.6). Therefore, we are recommending changes which would create a tighter bond between ID Protection and Sentinel, which would lower the number of false negatives in Sentinel. As a bonus, a security analyst would only need to work with a single dashboard to detect security threats related to the sign-in activity of user identities.

5.4 Sub-Research Question 3

In this section, we delve into the third research sub-question (1.3):

3. *How does the use of best-practice Conditional Access policies affect what is detected while using Microsoft's rule-sets from Content Hub and risk detection in ID Protection?*

As mentioned in section 2.4.3 we opted to follow Microsoft's recommendations for a Zero Trust architecture when selecting which CAPs to integrate. Now, the pivotal question arises: how does the use of best-practice Conditional Access policies affect what is detected? To address this we will use the results discovered in section 4.

Effectiveness of Conditional Access Policies

First and foremost it is important to remember that as mentioned in 2.4, CA acts as a gatekeeper for user authentication, controlling access rights. Out of the 14 policies that were enabled (3.5), only five policies were ever utilised in various combinations throughout the testing process, depending on the test:

- Require multifactor authentication for admins
- Require multifactor authentication for all users
- Require multifactor authentication for risky sign-ins
- No persistent browser session
- Require compliant or hybrid Azure AD joined device or multifactor authentication for all users

In addition to serving as gatekeepers for user authentication, CAPs play a crucial role as proactive safeguards against evolving security threats. By mandating MFA for various user categories and sign-in scenarios, these policies create

obstacles for potential malicious attackers, as seen during our testing process (4). Thereby significantly reducing the likelihood of unauthorised access (2.4.3). The implementation of session controls and device compliance checks further fortifies security measures, ensuring that only authenticated and authorised users gain access to sensitive resources. While Microsoft's recommendations provide a solid foundation for a Zero Trust architecture, organisations may need to customise these policies to align with their unique security requirements and regulatory obligations.

Moreover, it's essential to consider the effect of CAPs on what is detected, particularly concerning incidents and risk detections. During our testing (4), we observed that enabling CAPs influenced the types of incidents detected in Sentinel and risk detections in ID Protection. The results revealed that while most incidents remained consistent, there were variations in the incidents triggered and risk detection outcomes, indicating the impact of CAPs on the detection process.

Impact on Sentinel

The CAPs seamlessly integrate with Microsoft's rule sets from Content Hub, enhancing the overall security posture of our setup. The policies complement the already existing rule sets by adding a layer of access control. The integration ensures that security policies are consistently enforced across various Microsoft services and applications, thereby minimising the gaps and inconsistencies in our security framework.

When conducting the *Valid Account: New Location* attack (3.2.1), Sentinel created incidents triggered by the analytics rule *MFA explicit deny*, indicating attempts to sign in from new locations (4.2). However, upon enabling CAPs, the detection outcomes varied. While ID Protection detected the attack as before, Sentinel registered additional incidents that were flagged by analytics rules such as *Successful login from IP and failure from a different IP* and *Correlate Unfamiliar sign-in properties & atypical travel alerts*, showcasing the influence of CAPs on Sentinel's ability to detect possible attacks. Even though both test cases had MFA enabled, once CA was enabled the results correlated closer with what we were trying to test.

Impact on ID Protection

The implementation of CAPs has positively impacted risk detection capabilities provided by ID Protection. By enforcing best-practice policies, we have improved our ability to prevent security risks effectively. During our testing in section 4, we observed that some of the risk detections were automatically mitigated due to the conditions specified in the CAPs as seen in 4.5.2. This proactive approach to risk management, coupled with the strict access controls enforced by CA, has contributed to enhanced protection against potential threats.

Another noteworthy observation pertains to the impact CAPs had on risk detection, particularly in the context of the *Steal Web Session Cookie* attack scenario

(3.2.1). With CAPs enabled, a new risk detection with the detection type *Anomalous token* was triggered in ID Protection. It was registered with the attack type *Access using a valid account (Detected Offline)*, *Steal Web Session Cookie/Token Theft*, signalling a potential security threat related to cookie theft or session hijacking (4.7.2). This outcome underscores the value of CAPs in augmenting risk detection capabilities and fortifying the organisation's defence against cyber threats.

Conclusion

The biggest takeaway from testing with CA enabled was the fact that two new incidents were triggered in Sentinel and one new risk detection in ID Protection. Even though most of the results from simulating the attacks are the same whether CA has been enabled or not, some cases differed. The observed differences in the incident and risk detection outcomes underscore that there is a real benefit of implementing best-practice CAPs.

In conclusion, the addition of best-practice CAPs plays an important role in fortifying an organisation's security posture. We have seen that there is an increase of one more detection by both Sentinel and ID Protection with CA enabled, compared to without. Beyond just access control, CAPs serve as proactive deterrents against emerging threats, enhancing incident management capabilities, and facilitating seamless integration with other security solutions.

5.5 Sub-Research question 4

As mentioned in section 2.3.4, Microsoft asserts that only four specific risk detections can be triggered manually due to implemented machine learning, with three targeting user identities and one targeting workload identities. While simulating the different attacks, we kept these specific risk detections in mind to evaluate the quality of our testing. This gave way to our fourth research sub-question:

4. *Can we trigger any additional risk detections for user identities beyond those presented by Microsoft?*

When performing the different attacks, we aimed to trigger the three risk detections³ that Microsoft claims can be triggered manually for user identities. At the end of our testing procedure, we noticed we had managed to trigger all of these risk detections. There also exists a risk detection for *Leaked credentials*, for user identities, that was not triggered. Since Microsoft does not provide any guidance on how to simulate this risk detection for user identities, we chose to use the method outlined for workload identities, as mentioned in section 3.2.1. Still, it is somewhat unexpected that this went unnoticed by ID Protection, suggesting potential differences in the attack logic between user and workload identities.

However, when performing the attack *Steal Cookie*, we obtained an unexpected result, as presented in the result 4.7.2. Here we triggered the risk detection

³Anonymous IP address, Unfamiliar sign-in properties and Atypical travel, see section 2.3.4.

Anomalous token. This risk detection was not in the provided guide from Microsoft for simulating risk detections (2.3.4). Being able to trigger this risk detection, as well as the three others, indicates that we have been thorough with our testing methodology.

The goal of this research sub-question is for us to be able to assess the comprehensiveness of our testing method. As presented here all three risk detections related to user identities were triggered, and even one more. This demonstrates that our testing has been comprehensive, which leads to an increase in the confidence we can have in the reliability and credibility of our results.

5.6 Main Research Question

In this section, we will delve into the main research question (1.3):

Main research question: *How well do the rulesets provided by Microsoft in Sentinel Content Hub for Entra ID and Entra ID Protection with a best-practice setup of Conditional Access policies secure an organisation against user identity-based threats?*

To answer the main research question, we will utilise our findings from sub-research questions 1-4, discussed earlier in this chapter. Together, these four questions provide a broad spectrum of data for our assessment of the main research question.

The main research question aims to evaluate how well Sentinel works compared to ID Protection, with CA enabled. We do this to assess the importance of the three tools, and how they may be used individually or in combination to secure an organisation in the best possible way.

Attack Detection

Of the 15 attacks we performed, nine attacks were detected overall. This indicates that our chosen tools did discover 60% of the performed attacks, meaning that 40% of the attacks were undetected by our tools. There can be many reasons for this. One possible explanation is the utilisation of machine learning in both ID Protection and UEBA, as well as the specificity of the analytics rules in Sentinel. The use of machine learning might have impacted our results, leading to some alerts not being created. Microsoft does not provide any information about how their machine learning detection methods work. Therefore, we are not able to confirm whether any other alerts than those from our results would be present in a real attack scenario.

In any organisation, it would be better to have a higher percentage of detected attacks, as this implies fewer attacks slipping through undetected. Therefore, it is important to acknowledge that relying solely on one security tool would lower the detection percentage, given that in most cases, only one of the tools identified

the security incident. Consequently, we recommend having both Sentinel and ID Protection enabled and using them together effectively.

Sub-research question 1

When evaluating *how well* the rulesets in Sentinel secure an organisation about user identity-based threats, we chose to compare it to ID Protection. This is because it also performs detection of user identities within Azure and is embedded as a core part of ME-ID. As discussed in 5.2, there exist some differences between these two tools, but the most important difference for the main research question is the variance between the attack detection of these tools. Even though Sentinel detected a higher number of attacks compared to ID Protection, we would not recommend using only Sentinel as a detection system. The reason for this is that only *one* attack was discovered by both tools, meaning that the other detected attacks were detected by only one of the tools. This indicates that the tools add value to one another, rather than leading to an increased number of unnecessary detections. Again, since only *one* attack was triggered in both systems, we can state that the chance for alert fatigue (2.6) should not increase by having both tools in use simultaneously. Therefore, we do not see any reason for having only Sentinel enabled, since ID Protection adds value to the detection of attacks.

Sub-research question 2

When evaluating the effectiveness of Sentinel, we found it meaningful to consider the number of false positive and false negative results. As stated in 5.3.1, we did not encounter any false positives, indicating that the existing analytics rules do not require any changes. However, we did identify some false negatives (5.3.2), suggesting that additional analytics rules need to be implemented to cover a broader spectrum of attacks.

Our recommendation is to create additional analytics rules to detect and trigger an incident when a risk detection is made in ID Protection (5.3.3). This will enable organisations to continue using both Sentinel and ID Protection while eliminating the need for separate dashboards. This may reduce the risk of alert fatigue (2.6), as there would be no need to acknowledge alerts for the same incidents in both tools simultaneously. The benefit of this approach is a more secure system, leveraging the advantages of both Sentinel and ID Protection, presented in a single dashboard rather than two. Furthermore, it is possible to further use the SOAR features found in Sentinel to perform remediating actions in ID Protection (2.5.4), but this feature is not covered by the scope of this thesis (1.8).

Sub-Research question 3

When evaluating the effect of having the recommended CAPs enabled for the rulesets in Sentinel, we have chosen to investigate whether the number of detections changed. As stated in 5.4, the implementation of Microsoft's recommended CAPs

led to an increase in the detection of attacks, indicating that having CAPs enabled will better secure an organisation against user identity-based threats compared to not having them enabled.

Another important observation is the impact of CA when we performed the attacks. The MFA, that comes with having CA enabled, prevented access to our system in every scenario which had the goal of gaining access⁴, except for *Steal Web Session Cookie*. This indicates that having CA enabled effectively secures an organisation against threats.

However, even though it is beneficial that the attacks are being prevented by CA, it can be valuable to receive some form of alert or alarm to notify about a threat against your system, even if it is stopped. This is what we experienced Sentinel did, adding value to an organisation's knowledge about their threat landscape.

Sub-Research question 4

As shown in 5.5, we managed to trigger all the risk detections regarding user identities that Microsoft states are possible, and even one more. This demonstrates that our testing has been thorough and that our results have covered many different scenarios. As a result, our conclusion is based on a varied and broad foundation which gives us high confidence in our final evaluation of the main research question.

Conclusion

Our investigation of how well the rulesets provided by Microsoft in Content Hub for ME-ID and ID Protection, with a best-practice setup of CAPs, secure an organisation against user identity-based threats has provided us with valuable insights. By utilising findings from research sub-questions 1-4 (1.3), we have obtained a comprehensive understanding of how these tools perform in various attack scenarios.

Our main research question aimed to evaluate how well Sentinel works compared to ID Protection, with CA enabled, and to assess the importance of these tools individually or in combination for security in an organisation. Our findings indicate that while Sentinel and CA can secure an organisation well, using ID Protection in addition will increase the overall detection percentage. This suggests that using all three tools together will provide a more secure defence against user identity-based threats.

However, if an organisation determined that Sentinel would be their sole platform for detecting security incidents, our research has revealed areas needing improvement within the analytics rules provided by Microsoft. Our recommendations include creating analytics rules which integrate risk detections from the ID Protection dashboard into the incident dashboard in Sentinel for a unified monitoring solution. This solution could reduce alarm fatigue.

⁴These are the attack techniques under the tactics initial access and credential access 2.7.

The implementation of recommended CAPs led to an increase in the detection of attacks, showing the importance of having CA enabled. Additionally, we found value in receiving alerts in Sentinel, even for prevented attacks, to maintain awareness of potential threats.

Our testing, which successfully triggered one more risk detection than Microsoft provided a method for, gives us confidence in the accuracy and reliability of our results. Therefore, we are confident in the evaluation that the security measures offered to organisations by following Microsoft's recommendations for configuring CAPs and setting up Sentinel with Microsoft's rulesets for ME-ID and ID Protection from Content Hub, will secure organisations well against threats targeting user-identities, provided they also keep up with risk detections in the ID Protection dashboard. Overall, our evaluation highlights the significance of utilising a combination of Sentinel, ID Protection, and CA to secure organisations against user identity-based threats in the best possible way.

Chapter 6

Conclusion

The primary objective of this thesis was to explore Microsoft's setup of Sentinel and ID Protection, aiming to evaluate their efficiency in protecting organisations against user identity-based threats. To do this, we adopted a practical approach rooted in real-world attack simulations, referencing established frameworks such as MITRE ATT&CK and Microsoft's guide for simulating risk detections, and delved into the differences between these security solutions.

6.1 Answering the Problem Statement

In alignment with our problem statement (1.2), we investigated the effectiveness of the rulesets provided by Microsoft in Content Hub for ME-ID and ID Protection, particularly in detecting incoming attacks against users. Our approach aimed to simulate an organisation's security setup following Microsoft's recommendations by enabling the pre-made analytics rules in Content Hub and implementing best practice CAPs.

Furthermore, our testing methodology, rooted in real-world attack techniques, provided a robust framework for assessing the comprehensiveness of our evaluation. By referencing established frameworks and guidelines, we ensured a thorough and systematic assessment of the security setup, uncovering additional risk detections beyond the baseline provided by Microsoft's guide.

Our findings reveal that while both Sentinel and ID Protection offer distinct features and detection mechanisms, the combination of the two provides a broad and thorough approach to security. Notably, ID Protection focuses on providing detailed information about risk detections, while Sentinel offers a proactive incident management dashboard with a broader range of options for handling security incidents. Additionally, our testing revealed the significance of implementing best-practice CAPs, which serve as proactive measures against emerging threats and enhance incident management and detection capabilities. The addition of CAPs resulted in an increase in detections by both Sentinel and ID Protection, highlighting the importance of interlocked security measures. However, our research also identified areas for improvement in Microsoft's provided rulesets, emphasising the

importance of tighter integration between ID Protection and Sentinel to minimise false negatives and reduce alert fatigue.

6.2 Project Plan and Goals

During our preparatory work for this thesis, we constructed a plan for the project and the goals we wished to achieve. In this section, we will review how our work measured up to our initial expectations.

6.2.1 Effect and Project Goals

Throughout this thesis, we as a group aimed to achieve all of our effect goals (1.4) and project goals (1.5). When considering if we were able to achieve our goals, it is essential to consider the process as a whole, as the goals build upon the research questions and problem statement.

First, our thesis successfully addressed the effect goals (1.4) by evaluating the security measures implemented for user identities within an organisational setup following Microsoft's recommendations. This evaluation was carried out through a practical approach involving rigorous testing, demonstrating the effectiveness of the security measures implemented and thereby achieving our effect goals.

Moving on to the project goals, project goals 1 and 2 (1.5) were addressed throughout the testing phase of the thesis, as seen in section 3.2.1 and our results in section 4. Testing the rules provided by Content Hub also introduced us to what limits and difficulties are present when trying to simulate attacks. Additionally, we collected and examined the impact of using best-practice CAPs on incident detection in Sentinel, thus fulfilling project goals 3 and 4. Upon reviewing the whole process, it becomes evident that project goal 5 was met, as we demonstrated the differences between using the Sentinel dashboard versus the ID Protection dashboard.

Achieving the project and effect goals boosts team morale and highlights our commitment to clear objectives and dedicated work. This experience has not only deepened our understanding of the subject but also reinforced our belief in creating a cohesive task that encompasses a diverse range of activities, contributing to a rich and comprehensive learning experience.

6.2.2 Gantt Chart

At the beginning of this project, we planned our work for the semester by creating a Gantt chart (see D). Comparing the chart and the group's progress over the semester, we see that we closely followed our initial plan, ensuring alignment with our project objectives. Despite encountering a slight delay in starting the testing process, we successfully compensated for lost time and regained momentum. This ability to adapt shows our team's commitment to meeting deadlines. Overall, we maintained a positive trajectory, completing the initial draft by the scheduled date,

although the final submission was delayed by a few days. This journey highlights our capacity for evaluation and adjustment, ensuring progress despite challenges.

6.3 Further Research

Before delving into potential areas for further research and improvement, it's important to acknowledge the constraints and limitations of our current study. Engaging in a project of this scale is, to begin with, restricted by factors such as scope and available resources. While there are several areas for exploration, certain endeavours are outside the scope of this project. Expanding the scope too extensively could have required an entirely separate thesis to thoroughly explore certain areas.

With these considerations in mind, we as a group acknowledge that there are several areas where we could have made improvements or taken different approaches. Engaging in a project of this scale is always a learning process, and reflecting on our experiences helps refine our methods and insights.

6.3.1 Areas for Improvement

Looking ahead, several aspects could be explored further. Delving into aspects such as refining detection mechanisms and understanding security vulnerabilities beyond user identities could have been a valuable addition. We could have explored the nuances of false positives and their role in providing valuable insights into refining mechanisms and minimising alert fatigue. By investigating how the Sentinel solution provided by Microsoft tackled false positives more closely, we could have gained insights into refining mechanisms and minimizing alert fatigue, warranting further investigation.

Another path we could have taken involves exploring deeper into Sentinels SOAR functionality, which could uncover opportunities for improving incident detection and response workflows, ultimately enhancing efficiency. Expanding our scope to include other Microsoft security solutions, such as Azure Security Center and Microsoft Defender for Identity, could also provide a more comprehensive understanding of Microsoft's overall security posture. Comparing and contrasting the capabilities of all these solutions with Sentinel and ID Protection could offer valuable insights into how organisations can leverage Microsoft's suite of security tools to enhance their defences.

In addition to these aspects, we could have explored further, establishing an Office environment for testing purposes and further examination could have provided valuable insights.

6.3.2 Conclusion and Future Directions

As a final note, our research has provided valuable insights into the effectiveness of Sentinel and ID Protection in protecting organisations against user identity-based

threats. Based on our research, the combination of Sentinel and ID Protection offers a broad and thorough approach to security, with the addition of CAPs proving to be particularly significant. Nevertheless, there remain several opportunities for further research and exploration. By continuing to refine these methodologies and expanding the scope of inquiry, future work and research can contribute to the ongoing evolution of cyber security practices and technologies.

Bibliography

- [1] Microsoft, *What are microsoft entra sign-in logs?* <https://learn.microsoft.com/en-us/entra/identity/monitoring-health/concept-sign-ins>, Accessed on 01-05-2024, 2024.
- [2] American Registry for Internet Numbers (ARIN), *Autonomous system numbers*, <https://www.arin.net/resources/guide/asn/>, Accessed on 13-03-2024, 2024.
- [3] Microsoft, *What are the microsoft entra user provisioning logs?* <https://learn.microsoft.com/en-us/entra/identity/monitoring-health/concept-provisioning-logs>, Accessed on 01-05-2024, 2024.
- [4] Microsoft, *Advanced threat detection with user and entity behavior analytics (ueba) in microsoft sentinel*, <https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics>, Accessed on 16-05-2024, 2024.
- [5] Statistisk Sentralbyrå (SSB), *10966: Buy cloud computing services (per cent), by industry (sic2007), contents, year and employed*, <https://www.ssb.no/en/statbank/table/10966/tableViewLayout1/>, Accessed on 26-02-2024, 2023.
- [6] Radiant Logic, *New study reveals identity sprawl plagues organizations*, <https://www.radiantlogic.com/news/new-study-reveals-identity-sprawl-plagues-organizations-with-60-percent-reporting-over-21-disparate-identities-per-user/>, Accessed on 26-02-2024, 2022.
- [7] Politiets Sikkerhetstjeneste (PST), *Norsk trusselvurdering*, https://www.pst.no/globalassets/2023/ntv/ntv_2023_nor_web.pdf, 2023.
- [8] S. Hofmann, *Ny ibm-rapport: Kostnaden av et databrudd i 2023*, <https://www.cyberpilot.io/no/cyberpilot-blog/ny-ibm-rapport-kostnaden-av-et-databrudd>, Accessed on 08-03-2024, 2023.
- [9] EY, *Ey norwegian cloud maturity survey 2019*, https://www.digi.no/filer/B19001no-Norwegian_Cloud_Maturity_Survey_08__1_{}.pdf, 2019.
- [10] P.Krishna, *Introducing microsoft sentinel content hub!* <https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/introducing-microsoft-sentinel-content-hub/ba-p/2928102>, Accessed on 26-02-2024, 2021.

- [11] Microsoft, *What is identity and access management (iam)?* <https://www.microsoft.com/en-us/security/business/security-101/what-is-identity-access-management-iam>, Accessed on 08-03-2024.
- [12] Microsoft, *What is identity and access management (iam)?* <https://learn.microsoft.com/en-us/entra/fundamentals/introduction-identity-access-management>, Accessed on 08-03-2024, 2024.
- [13] Microsoft, *Identity and access management (iam) fundamental concepts*, <https://learn.microsoft.com/en-us/entra/fundamentals/identity-fundamental-concepts>, Accessed on 18-03-2024, 2024.
- [14] Microsoft, *Azure active directory is now microsoft entra id*, <https://www.microsoft.com/en-us/security/business/identity-access/microsoft-entra-id>, Accessed on 13-03-2024.
- [15] Microsoft, *New name for azure active directory*, <https://learn.microsoft.com/nb-no/entra/fundamentals/new-name>, Accessed on 13-03-2024, 2024.
- [16] Microsoft, *What is microsoft entra id?* <https://learn.microsoft.com/en-us/entra/fundamentals/whatis>, Accessed on 14-03-2024, 2023.
- [17] Microsoft, *What is conditional access?* <https://learn.microsoft.com/en-us/entra/identity/conditional-access/overview>, Accessed on 07-03-2024, 2023.
- [18] Microsoft, *Enterprise user management documentation*, <https://learn.microsoft.com/en-us/entra/identity/users/>, Accessed on 15-03-2024.
- [19] Microsoft, *What is hybrid identity with microsoft entra id?* <https://learn.microsoft.com/en-us/entra/identity/hybrid/whatis-hybrid-identity>, Accessed on 15-03-2024, 2023.
- [20] Microsoft, *What is identity protection?* <https://learn.microsoft.com/en-us/entra/id-protection/concept-identity-protection-risks>, Accessed on 07-03-2024, 2024.
- [21] Microsoft, *What is microsoft entra monitoring and health?* <https://learn.microsoft.com/en-us/entra/identity/monitoring-health/overview-monitoring-health>, Accessed on 15-03-2024, 2023.
- [22] Microsoft, *How to create, invite, and delete users*, <https://learn.microsoft.com/en-us/entra/fundamentals/how-to-create-delete-users>, Accessed on 18-03-2024, 2024.
- [23] Microsoft, *Groups in microsoft 365 and azure, and which is right for you*, <https://learn.microsoft.com/en-us/microsoft-365/community/all-about-groups?source=recommendations>, Accessed on 15-03-2024, 2023.

- [24] Microsoft, *Learn about groups and access rights in microsoft entra id*, <https://learn.microsoft.com/en-us/entra/fundamentals/concept-learn-about-groups>, Accessed on 18-03-2024, 2024.
- [25] Microsoft, *Prevent attacks using smart lockout - microsoft entra id*, <https://learn.microsoft.com/en-us/entra/identity/authentication/howto-password-smart-lockout>, Accessed on 13-04-2024, 2023.
- [26] Microsoft, *Microsoft entra id protection dashboard (preview)*, <https://learn.microsoft.com/en-us/entra/id-protection/id-protection-dashboard>, Accessed on 12-03-2024, 2023.
- [27] Microsoft, *What are risk detections?* <https://learn.microsoft.com/en-us/entra/id-protection/concept-identity-protection-risks>, Accessed on 07-03-2024, 2024.
- [28] The Tor Project, *Tor project | anonymity online*, <https://www.torproject.org/>, Accessed on 14-03-2024.
- [29] Microsoft, *Securing workload identities*, <https://learn.microsoft.com/en-us/entra/id-protection/concept-workload-identity-risk>, Accessed on 12-03-2024, 2024.
- [30] MITRE, *Valid accounts*, <https://attack.mitre.org/techniques/T1078/>, Accessed on 14-03-2024, 2023.
- [31] MITRE, *Account manipulation*, <https://attack.mitre.org/techniques/T1098/>, Accessed on 14-03-2024, 2023.
- [32] MITRE, *Brute force: Password spraying*, <https://attack.mitre.org/techniques/T1110/003/>, Accessed on 14-03-2024, 2023.
- [33] MITRE, *Gather victim identity information: Credentials*, <https://attack.mitre.org/techniques/T1589/001/>, Accessed on 14-03-2024, 2023.
- [34] MITRE, *Proxy*, <https://attack.mitre.org/techniques/T1090/>, Accessed on 14-03-2024, 2021.
- [35] MITRE, *Steal web session cookie*, <https://attack.mitre.org/techniques/T1539/>, Accessed on 14-03-2024, 2023.
- [36] Microsoft, *Simulating risk detections in microsoft entra id protection*, <https://learn.microsoft.com/en-us/entra/id-protection/howto-identity-protection-simulate-risk>, Accessed on 15-03-2024, 2024.
- [37] Broadcom, *What is a virtual machine?* <https://www.vmware.com/topics/glossary/content/virtual-machine.html.html>, Accessed on 15-03-2024.
- [38] MDN, *User agent*, https://developer.mozilla.org/en-US/docs/Glossary/User_agent, Accessed on 15-03-2024, 2023.
- [39] GitHub, *The tools you need to build what you want*. <https://github.com/features>, Accessed on 15-03-2024, 2022.

- [40] Microsoft, *Plan a conditional access deployment*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/plan-conditional-access#recommendations>, Accessed on 07-03-2024, 2024.
- [41] Microsoft, *Conditional access design principles and dependencies*, <https://learn.microsoft.com/en-us/azure/architecture/guide/security/conditional-access-design>, Accessed on 07-03-2024, 2024.
- [42] Microsoft, *What is conditional access report-only mode?* <https://learn.microsoft.com/en-us/entra/identity/conditional-access/concept-conditional-access-report-only>, Accessed on 01-04-2024, 2024.
- [43] Microsoft, *Risk-based access policies*, <https://learn.microsoft.com/en-us/entra/id-protection/concept-identity-protection-policies>, Accessed on 08-03-2024, 2024.
- [44] Microsoft, *What is zero trust?* <https://learn.microsoft.com/en-us/security/zero-trust/zero-trust-overview>, Accessed on 12-03-2024, 2024.
- [45] Microsoft, *Conditional access for zero trust*, <https://learn.microsoft.com/en-us/azure/architecture/guide/security/conditional-access-zero-trust>, Accessed on 12-03-2024, 2023.
- [46] Microsoft, *Conditional access architecture and personas*, <https://learn.microsoft.com/en-us/azure/architecture/guide/security/conditional-access-architecture>, Accessed on 13-03-2024.
- [47] Microsoft, *Conditional access templates*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/concept-conditional-access-policy-common?tabs=zero-trust>, Accessed on 08-03-2024, 2023.
- [48] Microsoft, *Common conditional access policy: Require mfa for administrators*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/howto-conditional-access-policy-admin-mfa>, Accessed on 08-03-2024, 2023.
- [49] Microsoft, *Common conditional access policy: Securing security info registration*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/howto-conditional-access-policy-registration>, Accessed on 13-03-2024, 2023.
- [50] Microsoft, *Common conditional access policy: Block legacy authentication*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/howto-conditional-access-policy-block-legacy>, Accessed on 08-03-2024, 2024.
- [51] Microsoft, *Tutorial: Enforce multifactor authentication for b2b guest users*, <https://learn.microsoft.com/en-us/entra/external-id/b2b-tutorial-require-mfa>, Accessed on 13-03-2024, 2024.

- [52] Microsoft, *Common conditional access policy: Require mfa for azure management*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/howto-conditional-access-policy-azure-management>, Accessed on 08-03-2024, 2023.
- [53] Microsoft, *Common conditional access policy: Sign-in risk-based multifactor authentication*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/howto-conditional-access-policy-risk>, Accessed on 08-03-2024, 2023.
- [54] Microsoft, *Common conditional access policy: Require reauthentication and disable browser persistence*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/howto-policy-persistent-browser-session>, Accessed on 13-03-2024, 2023.
- [55] Tenable, *1.1.3 ensure sign-in frequency is enabled and browser sessions are not persistent for administrative users*, https://www.tenable.com/audits/items/CIS_Microsoft_365_v2.0.0_E3_Level_1.audit:7d89f960b1c74ead376ecb9de44d23c7, Accessed on 13-03-2024.
- [56] Microsoft, *Common conditional access policy: Require approved client apps or app protection policy*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/howto-policy-approved-app-or-app-protection>, Accessed on 13-03-2024, 2023.
- [57] Microsoft, *Common conditional access policy: Require a compliant device, microsoft entra hybrid joined device, or multifactor authentication for all users*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/howto-conditional-access-policy-compliant-device>, Accessed on 13-03-2024, 2024.
- [58] Microsoft, *Use compliance policies to set rules for devices you manage with intune*, <https://learn.microsoft.com/en-us/mem/intune/protect/device-compliance-get-started>, Accessed on 13-03-2024, 2023.
- [59] O. 3. Reports, *Require mfa for microsoft 365 admin portals with conditional access*, <https://o365reports.com/2023/11/03/require-mfa-for-microsoft-365-admin-portals-with-conditional-access/>, Accessed on 13-03-2024, 2024.
- [60] Microsoft, *Common conditional access policy: Require multifactor authentication for admins accessing microsoft admin portals*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/how-to-policy-mfa-admin-portals>, Accessed on 13-03-2024, 2023.
- [61] Microsoft, *Common conditional access policy: Block access for users with insider risk (preview)*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/how-to-policy-insider-risk>, Accessed on 01-04-2024, 2024.

- [62] Microsoft, *Conditional access: Conditions*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/concept-conditional-access-conditions#insider-risk-preview>, Accessed on 01-04-2024, 2024.
- [63] Microsoft, *Learn about microsoft purview*, <https://learn.microsoft.com/en-us/purview/purview>, Accessed on 01-04-2024, 2024.
- [64] Microsoft, *Help dynamically mitigate risks with adaptive protection (preview)*, <https://learn.microsoft.com/en-us/purview/insider-risk-management-adaptive-protection?tabs=purview-portal>, Accessed on 01-04-2024, 2024.
- [65] Microsoft, *Microsoft sentinel*, <https://azure.microsoft.com/en-us/products/microsoft-sentinel>, Accessed on 28-02-2024, 2024.
- [66] Microsoft, *What is siem?* <https://www.microsoft.com/en-us/security/business/security-101/what-is-siem>, Accessed on 07-03-2024, 2024.
- [67] IBM, *What is security information and event management (siem)?* <https://www.ibm.com/topics/siem>, Accessed on 07-03-2024, 2024.
- [68] IBM, *What is soar?* <https://www.ibm.com/topics/security-orchestration-automation-response>, Accessed on 07-03-2024, 2024.
- [69] Microsoft, *What is soar?* <https://www.microsoft.com/en-us/security/business/security-101/what-is-soar>, Accessed on 07-03-2024, 2024.
- [70] Microsoft, *What is microsoft sentinel?* <https://learn.microsoft.com/en-us/azure/sentinel/overview>, Accessed on 08-03-2024, 2023.
- [71] Microsoft, *Microsoft sentinel data connectors*, <https://learn.microsoft.com/en-us/azure/sentinel/connect-data-sources>, Accessed on 08-03-2024, 2023.
- [72] Microsoft, *Visualize and monitor your data by using workbooks in microsoft sentinel*, <https://learn.microsoft.com/en-us/azure/sentinel/monitor-your-data>, Accessed on 08-03-2024, 2023.
- [73] Microsoft, *Create custom analytics rules to detect threats*, https://learn.microsoft.com/nb-no/azure/sentinel/detect-threats-custom?WT.mc_id=azuresentinel_portalcard_inproduct_analytics, Accessed on 08-03-2024, 2023.
- [74] Microsoft, *Automate threat response in microsoft sentinel with automation rules*, <https://learn.microsoft.com/en-us/azure/sentinel/automate-incident-handling-with-automation-rules>, Accessed on 08-03-2024, 2022.
- [75] Microsoft, *Tutorial: Respond to threats by using playbooks with automation rules in microsoft sentinel*, <https://learn.microsoft.com/en-us/azure/sentinel/tutorial-respond-threats-playbook?tabs=LAC%2Cincidents>, Accessed on 08-03-2024, 2023.

- [76] Microsoft, *Microsoft sentinel out-of-the-box content centralization changes*, <https://learn.microsoft.com/en-us/azure/sentinel/sentinel-content-centralize>, Accessed on 12-03-2024, 2024.
- [77] Microsoft, *About microsoft sentinel content and solutions*, <https://learn.microsoft.com/en-us/azure/sentinel/sentinel-solutions>, Accessed on 12-03-2024, 2023.
- [78] Microsoft, *Microsoft entra id solution for sentinel*, <https://azuremarketplace.microsoft.com/en/marketplace/apps/azuresentinel.azure-sentinel-solution-azureactivedirectory?tab=Overview>, Accessed on 12-03-2024.
- [79] Microsoft, *Microsoft entra id protection*, <https://azuremarketplace.microsoft.com/en/marketplace/apps/azuresentinel.azure-sentinel-solution-azureactivedirectoryip?tab=Overview>, Accessed on 12-03-2024.
- [80] N. Nursee, *Deep dive into microsoft sentinel's new overview dashboard*, <https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/deep-dive-into-microsoft-sentinel-s-new-overview-dashboard/ba-p/3860688>, Accessed on 30-04-2024, 2024.
- [81] Malwarebytes, *What is alert fatigue?* <https://www.threatdown.com/glossary/what-is-alert-fatigue/>, Accessed on 13-04-2024.
- [82] Merriam-Webster, *Definition of false negative*, <https://www.merriam-webster.com/dictionary/false+negative>, Accessed on 13-04-2024, 2024.
- [83] MITRE, *Mitre att&ck*, <https://attack.mitre.org/>, Accessed on 14-03-2024.
- [84] MITRE, *Reconnaissance*, <https://attack.mitre.org/tactics/TA0043/>, Accessed on 24-03-2024, 2020.
- [85] MITRE, *Initial access*, <https://attack.mitre.org/tactics/TA0001/>, Accessed on 27-03-2024, 2019.
- [86] MITRE, *Defense evasion*, <https://attack.mitre.org/tactics/TA0005/>, Accessed on 28-03-2024, 2019.
- [87] MITRE, *Credential access*, <https://attack.mitre.org/tactics/TA0006/>, Accessed on 27-03-2024, 2019.
- [88] MITRE, *Multi-factor authentication request generation*, <https://attack.mitre.org/techniques/T1621/>, Accessed on 27-03-2024, 2023.
- [89] A. Singh, *How and where are issued tokens saved within azuread*, <https://learn.microsoft.com/en-us/answers/questions/491230/how-and-where-are-issued-tokens-saved-within-azure>, Accessed on 28-03-2024, 2021.

- [90] Microsoft, *Web browser cookies used in microsoft entra authentication - microsoft entra id*, <https://learn.microsoft.com/en-us/entra/identity/authentication/concept-authentication-web-browser-cookies>, Accessed on 01-05-2024, 2023.
- [91] Microsoft, *Bypassing mfa with the pass-the-cookie attack*, <https://blog.netwrix.com/2022/11/29/bypassing-mfa-with-pass-the-cookie-attack/>, Accessed on 01-05-2024, 2022.
- [92] MITRE, *Brute force*, <https://attack.mitre.org/techniques/T1110/>, Accessed on 27-03-2024, 2023.
- [93] L. Grigas, *What is password hashing?* <https://nordpass.com/blog/password-hash/>, Accessed on 27-03-2024, 2023.
- [94] MITRE, *Brute force: Password guessing*, <https://attack.mitre.org/techniques/T1110/001>, Accessed on 27-03-2024, 2023.
- [95] MITRE, *Privilege escalation*, <https://attack.mitre.org/tactics/TA0004/>, Accessed on 28-03-2024, 2021.
- [96] MITRE, *Account manipulation: Additional cloud roles*, <https://attack.mitre.org/techniques/T1098/003/>, Accessed on 28-03-2024, 2023.
- [97] MITRE, *Command and control*, <https://attack.mitre.org/tactics/TA0011/>, Accessed on 28-03-2024, 2019.
- [98] MITRE, *Persistence*, <https://attack.mitre.org/tactics/TA0003/>, Accessed on 28-03-2024, 2019.
- [99] MITRE, *Create account*, <https://attack.mitre.org/techniques/T1136/>, Accessed on 28-03-2024, 2023.
- [100] MITRE, *Account manipulation: Additional cloud credentials*, <https://attack.mitre.org/techniques/T1098/001/>, Accessed on 28-03-2024, 2023.
- [101] Atomic Red Team, *Atomic test #3 - brute force credentials of single azure ad user*, <https://github.com/redcanaryco/atomic-red-team/blob/master/atomics/T1110.001/T1110.001.md>, Accessed on 05-04-2024.
- [102] Atomic Red Team, *Atomic test #4 - password spray all azure ad users with a single password*, <https://github.com/redcanaryco/atomic-red-team/blob/master/atomics/T1110.003/T1110.003.md>, Accessed on 05-04-2024.
- [103] Atomic Red Team, *Atomic red team*, <https://github.com/redcanaryco/atomic-red-team/blob/master/README.md>, Accessed on 05-04-2024.
- [104] C. Gagnier, <https://chromewebstore.google.com/detail/cookie-editor/hlkenndednhfkekhgdcidcfdnkalm>, Accessed on 30-03-2024.
- [105] Microsoft, *Enable per-user microsoft entra multifactor authentication to secure sign-in events*, <https://learn.microsoft.com/en-us/entra/identity/authentication/howto-mfa-userstates>, Accessed on 14-05-2024, 2023.

- [106] Microsoft, *Microsoft entra recommendation: Switch from per-user mfa to conditional access mfa*, <https://learn.microsoft.com/en-us/entra/identity/monitoring-health/recommendation-turn-off-per-user-mfa>, Accessed on 14-05-2024, 2024.
- [107] Microsoft, *Microsoft graph powershell overview*, <https://learn.microsoft.com/en-us/powershell/microsoftgraph/overview?view=graph-powershell-1.0>, Accessed on 01-04-2024, 2023.
- [108] Microsoft, *Conditional access templates*, <https://learn.microsoft.com/en-us/entra/identity/conditional-access/concept-conditional-access-policy-common?tabs=secure-foundation>, Accessed on 07-03-2024, 2023.
- [109] Microsoft, *How to troubleshoot microsoft entra sign-in errors*, <https://learn.microsoft.com/en-us/entra/identity/monitoring-health/howto-troubleshoot-sign-in-errors>, Accessed on 25-04-2024, 2024.
- [110] Microsoft, *Microsoft entra authentication & authorization error codes - microsoft identity platform*, <https://learn.microsoft.com/en-us/entra/identity-platform/reference-error-codes>, Accessed on 25-04-2024, 2024.
- [111] Microsoft, *Data connectors/template_azureactivedirectory.json*, https://github.com/Azure/Azure-Sentinel/blob/master/Solutions/Microsoft%20Entra%20ID/Data%20Connectors/template_AzureActiveDirectory.JSON, Accessed on 01-04-2024, 2023.
- [112] Microsoft, *Microsoft entra id/analytic rules*, <https://github.com/Azure/Azure-Sentinel/tree/master/Solutions/Microsoft%20Entra%20ID/Analytic%20Rules>, Accessed on 01-04-2024, 2024.
- [113] Microsoft, *Data connectors/template_azureactivedirectoryidentityprotection.json*, https://github.com/Azure/Azure-Sentinel/blob/master/Solutions/Microsoft%20Entra%20ID%20Protection/Data%20Connectors/template_AzureActiveDirectoryIdentityProtection.JSON, Accessed on 01-04-2024, 2023.
- [114] Microsoft, *Analytic rules/correlateipc_unfamiliar-atypical.yaml*, https://github.com/Azure/Azure-Sentinel/blob/master/Solutions/Microsoft%20Entra%20ID%20Protection/Analytic%20Rules/CorrelateIPC_Unfamiliar-Atypical.yaml, Accessed on 01-04-2024, 2024.

Appendix A

Additional Material

A.1 Script for Creating Users and Groups

This script is based on a script we made in the course "INFT2504 - Skytjenester som arbeidsflate" in our fifth semester. We then made some changes to it so that it would fit into our setup and the tests we were planning on doing.

```
#Installer MSGraph PowerShell v1.0
Install-Module Microsoft.Graph
Connect-MgGraph -Scopes "User.ReadWrite.All",
"Group.ReadWrite.All", "Directory.ReadWrite.All"

#CSV-fil til brukerne
$users = Import-Csv -Path '/path/to/users.csv' -Delimiter ";"

# --- OPPRETT GRUPPER FOR AVDELINGENE ---

#Opprett grupper for hver av avdelingene
$departments = @(
"Anonymous IP address",
"Unfamiliar sign-in properties",
"Atypical travel",
"Leaked credentials",
"Rest"
)

$departmentsNickname = @(
"AnonymousIPAddress",
"UnfamiliarSign-inProperties",
"AtypicalTravel",
"LeakedCredentials",
"Rest"
)

#Oppretter security groups til hver av avdelingene
#(som her er hva som skal testes)
foreach ($department in $departments) {
    #Henter ut gruppene som allerede eksisterer
    $existingGroup = Get-MgGroup |
    Where-Object { $_.DisplayName -eq $department }

    #Sjekker om gruppene finnes
    if ($existingGroup) {
        Write-Host "Gruppe med navn '$department'
eksisterer allerede."
    }
}
```

```

}
#Hvis gruppene ikke finnes opprettes de
else {
    $departmentInfo = @{
        displayName      = $department
        description      = "$department - Test Group"
        securityEnabled  = $true #true for security groups
        groupTypes       = @()
        mailEnabled      = $false #false for security groups
        mailNickname     = $departmentsNickname
        [$departments.IndexOf($department)]
    }

    #hvis avdelingen eksisterer,
    #opprett avdelingen til en gruppe
    if ($department) {
        New-MgGroup @departmentInfo
        Write-Host "Gruppe '$department' er opprettet"
    }
    #hvis avdelingen ikke eksisterer
    else {
        Write-Host "Avdelingen '$department' ble ikke funnet."
    }
}
}

# for å sjekke om avdelingsgruppene har blitt laget:
foreach ($group in $departments) {
    Get-MgGroup | Where-Object { $_.DisplayName -eq $group }
}

# --- OPPRETT BUKERE ---
# Opprett brukerkonto og tildel gruppe
foreach ($user in $users) {

    # Setter et passord som nå brukes for alle ansatte
    $Password = 'Passord123'

    #Setter passordet, og at bruker må opprette
    #nytt passord første gang de logger inn
    $PasswordProfile = @{
        Password                        = $Password
        ForceChangePasswordNextSignIn = $true
    }
}

```



```
}

#Sjekker om bruker finnes
if (-not (Get-MgUser -Filter "userPrincipalName eq
'$($user.UserPrincipalName)')) {
    #Oppretter de nye brukerne i en variabel
    #sånn at den kan brukes senere
    $newUser = New-MGUser
        -UserPrincipalName $user.UserPrincipalName `
        -DisplayName $user.DisplayName `
        -PasswordProfile $PasswordProfile `
        -AccountEnabled `
        -MailNickname $user.MailNickname `
        -Department $user.Department

    $groupName = $user.GroupName

    # Hent gruppen basert på GroupName
    # (viktig at groupname er det samme som
    # hvilken department de er i)
    $group = Get-MgGroup -Filter "displayName eq '$groupName'"
    "Hvis gruppen finnes:"
    if ($group) {
        # Tildel brukeren til gruppen ved
        # å bruke New-MgGroupMember
        foreach ($userToAdd in $newUser) {
            New-MgGroupMember -GroupId $group.Id
            -DirectoryObjectId $userToAdd.Id
            Write-Host "Bruker $($user.DisplayName)
                lagt til i gruppen med navn '$groupName'"
        }
    }
    #Gir beskjed dersom brukeren har en gruppe
    #som ikke eksisterer
    else {
        Write-Host "Gruppen '$groupName' ble ikke funnet."
    }
}

#Bruker finnes allerede
else {
    Write-Host "Bruker med UserPrincipalName
        $($user.UserPrincipalName) eksisterer allerede."
}
```

```
    }  
  }  
  
  # Sjekker antall medlemmer i hver av gruppene avdelingene  
  # + gruppa for alle ansatte  
  Write-Host "-- Antall medlemmer i gruppene --"  
  foreach ($group in ($departments)) {  
    $groups = Get-MgGroup |  
    Where-Object { $_.DisplayName -eq $group }  
  
    if ($groups) {  
      $members = Get-MgGroupMember -GroupId $groups.Id  
      $memberCount = $members.Count  
      Write-Host "Gruppe: $group, Antall medlemmer: $memberCount"  
    }  
    else {  
      Write-Host "Gruppa '$group' finnes ikke."  
    }  
  }  
}  
  
# --- OPPRETT FELLES GRUPPE FOR ALLE TEST-BRUKERE ---  
  
#Opprett en felles gruppe for alle ansatte  
#Lager en ny gruppe for alle hvis den ikke allerede eksisterer  
$allTestUsersGroup = Get-MgGroup |  
Where-Object { $_.DisplayName -eq "AllTestUsers" }  
  
if (-not $allTestUsersGroup) {  
  $allTestUsersInfo = @{  
    displayName      = "AllTestUsers"  
    description      = "Group containing all test-users"  
    securityEnabled  = $true  
    groupTypes       = @()  
    mailEnabled      = $false  
    mailNickname     = "alltestusers"  
  }  
  
  New-MgGroup @allTestUsersInfo  
  Write-Host "Gruppe med navn 'AllTestUsers' ble opprettet"  
}  
else {  
  Write-Host "Gruppe med navn 'AllTestUsers' eksisterer allerede."  
}
```

```
# Sjekker om gruppene for avdelingene
# og "allEmployees" har blitt laget
foreach ($group in ($departments + "AllTestUsers")) {
    Get-MgGroup | Where-Object { $_.DisplayName -eq $group }
}

# Legger til de ansatte fra de ulike
# avdelingene til gruppa "AllTestUsers"
#Merk: forutser at de ansatte har blitt opprettet
# og plassert i avdelingen/gruppa de tilhører
foreach ($department in $departments) {
    $departmentGroup = Get-MgGroup |
    Where-Object { $_.DisplayName -eq $department }

    #Hvis gruppa til avdelingene finnes så:
    if ($departmentGroup) {
        $members = Get-MgGroupMember -GroupId $departmentGroup.Id

        foreach ($member in $members) {
            # Sjekker om brukeren allerede ble lagt til i gruppa
            # Merk: Antar at når scriptet skal kjøres første gang
            # så er allEmployees gruppa helt tom
            # altså det er ingen i gruppa,
            # men for testing er dette greit å ha med
            $isMember = Get-MgGroupMember
            -GroupId $allTestUsersGroup.Id |
            Where-Object { $_.Id -eq $member.Id }

            if (-not $isMember) {
                # Legger brukere til "AllEmployees" gruppa
                New-MgGroupMember -GroupId $allTestUsersGroup.Id
                -DirectoryObjectId $member.Id

                $user = Get-MgUser -UserId $member.Id
                Write-Host " -- AllTestUsers Nye Medlemmer --"
                Write-Host "New member: $($user.DisplayName)"
            }
            else {
                $user = Get-MgUser -UserId $member.Id
                Write-Host "$($user.DisplayName) er allerede
                lagt til i gruppa AllTestUsers"
            }
        }
    }
}
```

```
    }  
  }  
  
  # Skriver ut alle medlemmene av "AllTestUsers" gruppa  
  $allTestUsersMembers = Get-MgGroupMember  
  -GroupId $allTestUsersGroup.Id  
  
  Write-Host "-- Medlemmer av gruppa 'AllTestUsers': --"  
  foreach ($member in $allTestUsersMembers) {  
    $user = Get-MgUser -UserId $member.Id  
    Write-Host "Navn: $($user.DisplayName),  
    UserPrincipalName: $($user.UserPrincipalName)"  
    # Prøvde å få den til å skrive ut avdelingene til de ansatte  
    # -> , "Avdeling: $($user.Department)",  
    # men det funka ikke helt (prøvde også med $($user.GroupName)  
  }  
  
  # Sjekker antall medlemmer i hver av gruppene avdelingene  
  # + gruppa for alle ansatte  
  Write-Host "-- Antall medlemmer i gruppene --"  
  foreach ($group in ($departments + "AllTestUsers")) {  
    $mgGroups = Get-MgGroup |  
    Where-Object { $_.DisplayName -eq $group }  
  
    if ($mgGroups) {  
      $members = Get-MgGroupMember -GroupId $mgGroups.Id  
      $memberCount = $members.Count  
      Write-Host "Gruppe: $group, Antall medlemmer: $memberCount"  
    }  
    else {  
      Write-Host "Gruppa '$group' finnes ikke."  
    }  
  }  
}  
  
# --- OPPDATER USAGE LOCATION PÅ ALLE BRUKERE I ALLTESTUSERS ---  
# Updating usage location on all test users  
  
# Get the group "AllTestUsers"  
$group = Get-MgGroup -DisplayName "AllTestUsers"  
  
if ($group) {  
  # Get all members of the group  
  $members = Get-MgGroupMember -GroupId $group.Id
```

```
if ($members) {
    # Update usage location for each member of the group
    foreach ($member in $members) {
        try {
            Update-MgUser -UserId $member.Id
            -UsageLocation "NO" -ErrorAction Stop
            Write-Host "Usage location updated for
                $($member.DisplayName)"
        } catch {
            Write-Host "Error updating usage location for
                $($member.DisplayName). Error: $_"
        }
    }
} else {
    Write-Host "No members found in the 'AllTestUsers' group."
}
} else {
    Write-Host "The group 'AllTestUsers' was not found."
}

#GIVE PERMISSION TO MANAGED IDENTITY
$managedIdentityObjectId = "efa58a0b-6a80-463c-8043-d61336c29d73"
$subscriptionId = "8fe266af-9a8d-40b0-bcb6-08d23e112c60"
#Assign 'User.Read.All' permission
New-AzRoleAssignment -ObjectId $managedIdentityObjectId
-RoleDefinitionName "User.Read.All"
-Scope "/subscriptions/{$subscriptionId}"

#Assign 'User.ReadWrite.All' permission
New-AzRoleAssignment -ObjectId $managedIdentityObjectId
-RoleDefinitionName "User.ReadWrite.All"
-Scope "/subscriptions/{$subscriptionId}"

#Assign 'Directory.Read.All' permission
New-AzRoleAssignment -ObjectId $managedIdentityObjectId
-RoleDefinitionName "Directory.Read.All"
-Scope "/subscriptions/{$subscriptionId}"

#Assign 'Directory.ReadWrite.All' permission
New-AzRoleAssignment -ObjectId $managedIdentityObjectId
-RoleDefinitionName "Directory.ReadWrite.All"
-Scope "/subscriptions/{$subscriptionId}"
```

A.2 Microsoft Entra ID - solution

Data Connectors

Microsoft Entra ID This data connector allows us to gain insight into ME-ID by connecting data logs to Microsoft Sentinel. The log types we have connected in our environment are sign-in logs, audit logs, non-interactive user sign-in logs (preview), user risk events (preview) and risky users (preview). We have used these logs to gather as much detailed information as possible about what the user identities do in our system, so we further can crosscheck what our system has captured regarding which security threats we initiated [111].

Analytics rules

MFA Rejected by User Gives us information when a user has rejected an MFA prompt. This could be harmless, but it can also be an indicator that a username and password have been compromised and a threat actor is trying to log into the account [112].

Attempt to bypass conditional access rule in Microsoft Entra ID This rule offers insight into attempts to bypass the configured CA rules. By detecting these attempts, valuable information is provided to optimize system security. These CA rules must be correctly configured to minimize loopholes and ensure effective enforcement [112].

Failed login attempts to Azure Portal This rule simply identifies many failed login attempts, or some failed login attempts from multiple IP addresses, into the Azure Portal, as this could indicate an attack, for instance a Bruteforce [112].

Account Created and Deleted in Short Timeframe This rule is used for finding accounts created and deleted in a timeframe of 24 hours. This could indicate that an attacker made a user for their use, and then deleted it when finished to remove suspicion [112].

Account created or deleted by non-approved user This rule identifies when user accounts are created or deleted by a non-approved user defined in a list [112].

Attempts to sign in to disabled accounts This rule identifies when there are multiple failed login attempts to disabled accounts [112].

MFA Spamming followed by Successful login This rule will identify MFA Spamming followed by a Successful login within a given time window of 5 minutes and with a Default Failure count of 10 [112].

Authentication Methods Changed for Privileged Account Identifies authentication methods being changed for a privileged account, as this could be an indication that an attacker is adding an authentication method to the privileged account for continued access [112].

- Suspicious Sign In Followed by MFA Modification** This rule uses Microsoft Sentinels UEBA features to look for suspicious logins followed by modifications to the MFA settings by that user [112].
- Successful logon from IP and failure from a different IP** This rule identifies when a user account successfully logs into an Azure App from one IP and within 10 minutes fails to log in to the same Azure App from another IP [112].
- Distributed Password cracking attempts in Microsoft Entra ID** This rule uses the Microsoft Entra ID SigninLogs to look for an unusually high amount of failed password attempts coming from multiple locations for a user account [112].
- Password spray attack against Microsoft Entra ID application** This rule will look for login failures from multiple accounts from the same IP address within a default time window of 20 minutes to identify possible password spray activity against Microsoft Entra ID applications [112].
- Brute force attack against Azure Portal** This rule detects Bruteforce attacks in the Azure Portal by monitoring multiple authentication failures, more than 10, following a successful login within a timeframe of 20 minutes [112].
- Multiple admin membership removals from newly created admin** This rule will detect when a newly created Global administrator is removing multiple existing global administrators [112].
- User Accounts - Sign in Failure due to CA Spikes** This rule will identify a spike¹ in failed logins from user accounts [112].
- Privileged Accounts - Sign in Failure Spikes** Same as the rule above, but this rule will help identify if it spikes in failed logins from privileged accounts [112].
- Sign-ins from IPs that attempt sign-ins to disabled accounts** This rule will identify instances where multiple IP addresses are attempting to sign in to one or more disabled accounts. These attempts occur from an IP address that has previously been used for successful sign-ins from other accounts [112].
- Explicit MFA Deny** This rule identifies when a user explicitly denies a MFA push/alert [112].
- New User Assigned to Privileged Role** Identifies if a privileged role is assigned to a new user, giving the user privileged access [112].
- User Assigned New Privileged Role** Identifies when a new eligible or active privileged role is assigned to a user, giving the user privileged access [112].

¹A sudden increase based on an event based on historical baseline values

Bulk Changes to Privileged Account Permission Identifies when changes to multiple users' permissions are changed at once, as this could enable an attacker's access to Azure subscriptions in an environment [112].

A.3 Microsoft Entra ID Protection - solution

Data connectors

Microsoft Entra ID Protection This data connector allows us to get a consolidated view of risky users and risky events, while also giving us the ability to remediate immediately by connecting playbooks and analytics rules [113].

Analytics rule

Correlate Unfamiliar sign-in properties and atypical travel alerts This analytic rule is enabled to give us an alert when the system gets the combination of the Unfamiliar sign-in properties alert and the Atypical travel alert within +10 minutes or -10 minutes window about the same user [114].

A.4 MFA explicitly deny, full table

| | |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TenantId | 3a69aa66-8e8b-4554-87cd-60ebebea42cd |
| SourceSystem | Azure AD |
| TimeGenerated [UTC] | 2024-04-09T07:51:43.4558273Z |
| ResourceId | /tenants/3b0e731d-dd91-4040-b0f4-3636e3bf415d/providers/Microsoft.aadiam |
| OperationName | Sign-in activity |
| OperationVersion | 1.0 |
| Category | SignInLogs |
| ResultType | 500121 |
| ResultSignature | None |
| ResultDescription | Authentication failed during strong authentication request. |
| DurationMs | 0 |
| Resource | Microsoft.aadiam |
| ResourceGroup | Microsoft.aadiam |
| Identity | Rest User 3 |
| Level | 4 |
| Location | NO |
| AppDisplayName | My Apps |
| AppId | 2793995e-0a7d-40d7-bd35-6968ba142197 |
| AuthenticationContextClassReferences | [{"id":"urn:user:registersecurityinfo","detail":"previouslySatisfied"}] |
| AuthenticationDetails | [{"authenticationStepDateTime":"2024-04-09T07:45:21.0629194+00:00","authenticationMethod":"Previously satisfied","succeeded":true,"authenticationStepResultDetail":"First factor requirement satisfied by claim in the token","authenticationStepRequirement":"Primary authentication","StatusSequence":0,"RequestSequence":0,"authenticationStepDateTime":"2024-04-09T07:45:41+00:00","authenticationMethod":"Mobile app notification","succeeded":false,"authenticationStepResultDetail":"Authentication denied; user declined the authentication","authenticationStepRequirement":"Primary authentication","StatusSequence":1712648741688,"RequestSequence":1712648741688}] |
| AuthenticationProcessingDetails | [{"key":"Legacy TLS (TLS 1.0, 1.1, 3DES)","value":"False","key":"Is CAE Token","value":"False"}] |
| AuthenticationRequirement | multiFactorAuthentication |
| AuthenticationRequirementPolicies | [{"requirementProvider":"user","detail":"Per-user MFA"}] |
| ClientAppUsed | Browser |
| ConditionalAccessStatus | notApplied |

Table A.1: Analytics Rule: MFA explicitly deny, part 1

| | |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CreatedDateTime [UTC] | 2024-04-09T07:45:21.0629194Z |
| DeviceDetail_dynamic | "deviceId":"","operatingSystem":"MacOs","browser":"Safari 16.4" |
| IsInteractive | true |
| Id | 02761c5f-9d4d-4407-bcf2-e51e20ed7c00 |
| IpAddress | 62.148.32.132 |
| LocationDetails_dynamic | "city":"Førnebu","state":"Akershus","countryOrRegion":"NO","geoC |
| MfaDetail_dynamic | "authMethod":"Mobile app notification" |
| OriginalRequestId | 02761c5f-9d4d-4407-bcf2-e51e20ed7c00 |
| ProcessingTimeInMilliseconds | 71 |
| RiskDetail | none |
| RiskLevelAggregated | none |
| RiskLevelDuringSignIn | none |
| ResourceDisplayName | Microsoft Graph |
| ResourceIdentity | 00000003-0000-0000-c000-000000000000 |
| ResourceServicePrincipalId | 7b8c4f71-4760-43c1-8c97-e58b1cee8051 |
| Status_dynamic | "errorCode":500121,"additionalDetails":"MFA denied; user declined the authentication","failureReason":"Authentication failed during strong authentication request." |
| TokenIssuerType | AzureAD |
| UserAgent | Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7) AppleWebKit/605.1.15 (KHTML, like Gecko) Version/16.4 Safari/605.1.15 |
| UserDisplayName | Rest User 3 |
| UserId | cd64db3f-2a67-40bc-a024-c82f09e7a09d |
| UserPrincipalName | restuser3@ntnuinf2504.onmicrosoft.com |
| AADTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| UserType | Member |
| ResourceTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| HomeTenantId | 3b0e731d-dd91-4040-b0f4-3636e3bf415d |
| UniqueTokenIdentifier | Xxx2Ak2dB0S88uUeIO18AA |
| AutonomousSystemNumber | 13243 |
| AuthenticationProtocol | none |
| CrossTenantAccessType | none |
| Type | SignInLogs |
| PublicIP | 62.148.32.132 |
| Name | restuser3 |
| UPNSuffix | ntnuinf2504.onmicrosoft.com |

Table A.2: Analytics Rule: MFA explicitly deny, part 2

A.5 Scripts used to perform the brute force attacks

A.5.1 Brute Force: Password Guessing

```

Import-Module -Name AzureAD

$passwords = cat ".\passwords.txt".split("`n")
$username = "victim2@ntnuinf2504.onmicrosoft.com"

foreach($password in $passwords) {
    $PWord = ConvertTo-SecureString -String "$password" -AsPlainText -Force
    $Credential = New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList $username, $PWord
    try {
        Write-Host " [-] Attempting ${password} on account ${username}."
        Connect-AzureAD -Credential $Credential # 2>&1> $null
        # if credentials aren't correct, it will break just above and goes into catch block, so if we're here we can display success
        Write-Host " [!] ${username}:${password} are valid credentials!`r`n"
        break
    } catch {
        <# if ($password -eq "Inft2504") {
            foreach ($err in $Error) {
                $lastErrorMessage = $err.Exception.Message
                if ($lastErrorMessage -like "*AADSTS*") {
                    Write-Host " [!!!] Error message: ${lastErrorMessage}.`r`n"
                }
            }
        } #>
        Write-Host " [-] ${username}:${password} invalid credentials.`r`n"
    }
}
Write-Host "End of bruteforce"

```

Figure A.1: Password Guessing Script

A.5.2 Brute Force: Password Spraying

```

Import-Module -Name AzureAD

$valid_password = "Toyota010234"
$Users = cat ".\users.txt".split("`n")

$PWord = ConvertTo-SecureString -String "${valid_password}" -AsPlainText -Force

foreach($user in $Users) {
    $Credential = New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList "$user", $PWord
    try {
        Write-Host " [-] Attempting ${valid_password} on account ${user}."
        Connect-AzureAD -Credential $Credential # 2>&1> $null
        # if credentials aren't correct, it will break just above and goes into catch block, so if we're here we can display success
        Write-Host " [!] ${user}:${valid_password} are valid credentials!`r`n"
        Disconnect-AzureAD > $null
    } catch {
        Write-Host " [-] ${user}:${valid_password} invalid credentials.`r`n"
    }
}
Write-Host "End of password spraying"

```

Figure A.2: Password Spraying Script

A.6 User CSV File

The table shows the content of our csv-file that contains all the users we created for testing. However, all the passwords to the users are removed from the file. Users that were created during testing are not included in this.

Table A.3: csv file with all the users created

| |
|---------------------------------------------------------------------------------------------------------------------------------------------------|
| UserPrincipalName;DisplayName;MailNickname;Department;Password;GroupName; |
| AnonUser1@ntnuinf2504.onmicrosoft.com;Anon User 1;anonuser1;Anonymous IP address;;Anonymous IP address; |
| AnonUser2@ntnuinf2504.onmicrosoft.com;Anon User 2;anonuser2;Anonymous IP address;;Anonymous IP address; |
| UnfamiliarSignin1@ntnuinf2504.onmicrosoft.com;Unfamiliar Signin 1;unfamiliarsignin1;Unfamiliar sign-in properties;;Unfamiliar sign-in properties; |
| UnfamiliarSignin2@ntnuinf2504.onmicrosoft.com;Unfamiliar Signin 2;unfamiliarsignin2;Unfamiliar sign-in properties;;Unfamiliar sign-in properties; |
| TravelUser1@ntnuinf2504.onmicrosoft.com;Travel User 1;traveluser1;Atypical travel;;Atypical travel; |
| TravelUser2@ntnuinf2504.onmicrosoft.com;Travel User 2;traveluser2;Atypical travel;;Atypical travel; |
| LeakedUser1@ntnuinf2504.onmicrosoft.com;Leaked User 1;leakeduser1;Leaked credentials;;Leaked credentials; |
| LeakedUser2@ntnuinf2504.onmicrosoft.com;Leaked User 2;leakeduser2;Leaked credentials;;Leaked credentials; |
| RestUser1@ntnuinf2504.onmicrosoft.com;Rest User 1;restuser1;Rest;;Rest; |
| RestUser2@ntnuinf2504.onmicrosoft.com;Rest User 2;restuser2;Rest;;Rest; |
| RestUser3@ntnuinf2504.onmicrosoft.com;Rest User 3;restuser3;Rest;;Rest; |
| RestUser4@ntnuinf2504.onmicrosoft.com;Rest User 4;restuser4;Rest;;Rest; |
| RestUser5@ntnuinf2504.onmicrosoft.com;Rest User 5;restuser5;Rest;;Rest; |
| RestUser6@ntnuinf2504.onmicrosoft.com;Rest User 6;restuser6;Rest;;Rest; |
| RestUser7@ntnuinf2504.onmicrosoft.com;Rest User 7;restuser7;Rest;;Rest; |
| RestUser8@ntnuinf2504.onmicrosoft.com;Rest User 8;restuser8;Rest;;Rest; |
| RestUser9@ntnuinf2504.onmicrosoft.com;Rest User 9;restuser9;Rest;;Rest; |
| RestUser10@ntnuinf2504.onmicrosoft.com;Rest User 10;restuser10;Rest;;Rest; |
| RestUser11@ntnuinf2504.onmicrosoft.com;Rest User 11;restuser11;Rest;;Rest; |
| RestUser12@ntnuinf2504.onmicrosoft.com;Rest User 12;restuser12;Rest;;Rest; |
| Victim1@ntnuinf2504.onmicrosoft.com;Victim 1;victim1;Unik users;;Unik user |
| Victim2@ntnuinf2504.onmicrosoft.com;Victim 2;victim2;Unik users;;Unik users |
| Victim3@ntnuinf2504.onmicrosoft.com;Victim 3;victim3;Unik users;;Unik users |
| Victim4@ntnuinf2504.onmicrosoft.com;Victim 4;victim4;Unik users;;Unik users |

Appendix B

Standard Agreement

Fastsatt av prorektor for utdanning 10.12.2020

STANDARDAVTALE

om utføring av studentoppgave i samarbeid med ekstern virksomhet

Avtalen er ufravikelig for studentoppgaver (heretter oppgave) ved NTNU som utføres i samarbeid med ekstern virksomhet.

Forklaring av begrep

Opphavsrett

Er den rett som den som skaper et åndsverk har til å fremstille eksemplarer av åndsverket og gjøre det tilgjengelig for allmennheten. Et åndsverk kan være et litterært, vitenskapelig eller kunstnerisk verk. En studentoppgave vil være et åndsverk.

Eiendomsrett til resultater

Betyr at den som eier resultatene bestemmer over disse. Utgangspunktet er at studenten eier resultatene fra sitt studentarbeid. Studenten kan også overføre eiendomsretten til den eksterne virksomheten.

Bruksrett til resultater

Den som eier resultatene kan gi andre en rett til å bruke resultatene, f.eks. at studenten gir NTNU og den eksterne virksomheten rett til å bruke resultatene fra studentoppgaven i deres virksomhet.

Prosjektbakgrunn

Det partene i avtalen har med seg inn i prosjektet, dvs. som vedkommende eier eller har rettigheter til fra før og som brukes i det videre arbeidet med studentoppgaven. Dette kan også være materiale som tredjepersoner (som ikke er part i avtalen) har rettigheter til.

Utsatt offentliggjøring

Betyr at oppgaven ikke blir tilgjengelig for allmennheten før etter en viss tid, f.eks. før etter tre år. Da vil det kun være veileder ved NTNU, sensorene og den eksterne virksomheten som har tilgang til studentarbeidet de tre første årene etter at studentarbeidet er innlevert.

1. Avtaleparter

| |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Norges teknisk-naturvitenskapelige universitet (NTNU) Institutt: Institutt for datateknologi og informatikk |
| Veileder ved NTNU: e-post og tlf.: |
| Ekstern virksomhet: Tietoevry Tech Services Ekstern virksomhet sin kontaktperson, e-post og tlf.: Thor Larsen, thor.larsen@tietoevry.com, mob: 90946017 |
| Student: Amund Fredrik Strømsnes Fødselsdato: 14.12.1999 |
| Student: Lea Arwen Utstøl Fødselsdato: 22.08.2002 |
| Student: Dina Hagen Steinskog Fødselsdato: 19.07.2002 |

Partene har ansvar for å klarere eventuelle immaterielle rettigheter som studenten, NTNU, den eksterne eller tredjeperson (som ikke er part i avtalen) har til prosjektbakgrunn før bruk i forbindelse med utførelse av oppgaven. Eierskap til prosjektbakgrunn skal fremgå av eget vedlegg til avtalen der dette kan ha betydning for utførelse av oppgaven.

2. Utførelse av oppgave

Studenten skal utføre: (sett kryss)

| | |
|-----------------|---|
| Masteroppgave | |
| Bacheloroppgave | X |
| Prosjektoppgave | |
| Annen oppgave | |

| |
|-----------------------|
| Startdato: 10.01.2024 |
| Sluttdato: 21.05.2024 |

Oppgavens arbeidstittel er:
How to monitor and respond to identity-based threats in Azure using Microsoft Entra ID Protection and Microsoft Sentinel

Ansvarlig veileder ved NTNU har det overordnede faglige ansvaret for utforming og godkjenning av prosjektbeskrivelse og studentens læring.

3. Ekstern virksomhet sine plikter

Ekstern virksomhet skal stille med en kontaktperson som har nødvendig faglig kompetanse til å gi studenten tilstrekkelig veiledning i samarbeid med veileder ved NTNU. Ekstern kontaktperson fremgår i punkt 1.

Formålet med oppgaven er studentarbeid. Oppgaven utføres som ledd i studiet. Studenten skal ikke motta lønn eller lignende godtgjørelse fra den eksterne for studentarbeidet. Utgifter knyttet til gjennomføring av oppgaven skal dekkes av den eksterne. Aktuelle utgifter kan for eksempel være reiser, materialer for bygging av prototyp, innkjøp av prøver, tester på lab, kjemikalier. Studenten skal klarere dekning av utgifter med ekstern virksomhet på forhånd.

| |
|----------------------------------------------------------------------------|
| Ekstern virksomhet skal dekke følgende utgifter til utførelse av oppgaven: |
|----------------------------------------------------------------------------|

Dekning av utgifter til annet enn det som er oppført her avgjøres av den eksterne underveis i arbeidet.

4. Studentens rettigheter

Studenten har opphavsrett til oppgaven¹. Alle resultater av oppgaven, skapt av studenten alene gjennom arbeidet med oppgaven, eies av studenten med de begrensninger som følger av punkt 5, 6 og 7 nedenfor. Eiendomsretten til resultatene overføres til ekstern virksomhet hvis punkt 5 b er avkrysset eller for tilfelle som i punkt 6 (overføring ved patenterbare oppfinnelser).

I henhold til lov om opphavsrett til åndsverk beholder alltid studenten de ideelle rettigheter til eget åndsverk, dvs. retten til navngivelse og vern mot krenkende bruk.

Studenten har rett til å inngå egen avtale med NTNU om publisering av sin oppgave i NTNUs institusjonelle arkiv på Internett (NTNU Open). Studenten har også rett til å publisere oppgaven eller deler av den i andre sammenhenger dersom det ikke i denne avtalen er avtalt begrensninger i adgangen til å publisere, jf. punkt 8.

5. Den eksterne virksomheten sine rettigheter

Der oppgaven bygger på, eller videreutvikler materiale og/eller metoder (prosjektbakgrunn) som eies av den eksterne, eies prosjektbakgrunnen fortsatt av den eksterne. Hvis studenten skal utnytte resultater som inkluderer den eksterne sin prosjektbakgrunn, forutsetter dette at det er inngått egen avtale om dette mellom studenten og den eksterne virksomheten.

Alternativ a) (sett kryss) Hovedregel

| | |
|-------------------------------------|------------------------------------------------------------------|
| <input checked="" type="checkbox"/> | Ekstern virksomhet skal ha bruksrett til resultatene av oppgaven |
|-------------------------------------|------------------------------------------------------------------|

¹ Jf. Lov om opphavsrett til åndsverk mv. av 15.06.2018 § 1

Dette innebærer at ekstern virksomhet skal ha rett til å benytte resultatene av oppgaven i egen virksomhet. Retten er ikke-eksklusiv.

Alternativ b) (sett kryss) Unntak

| | |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> | Ekstern virksomhet skal ha eiendomsretten til resultatene av oppgaven og studentens bidrag i ekstern virksomhet sitt prosjekt |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------|

Begrunnelse for at ekstern virksomhet har behov for å få overført eiendomsrett til resultatene:

6. Godtgjøring ved patenterbare oppfinnelser

Dersom studenten i forbindelse med utførelsen av oppgaven har nådd frem til en patenterbar oppfinnelse, enten alene eller sammen med andre, kan den eksterne kreve retten til oppfinnelsen overført til seg. Dette forutsetter at utnyttelsen av oppfinnelsen faller inn under den eksterne sitt virksomhetsområde. I så fall har studenten krav på rimelig godtgjøring. Godtgjøringen skal fastsettes i samsvar med arbeidstakeroppfinnelsesloven § 7. Fristbestemmelsene i § 7 gis tilsvarende anvendelse.

7. NTNU sine rettigheter

De innleverte filer av oppgaven med vedlegg, som er nødvendig for sensur og arkivering ved NTNU, tilhører NTNU. NTNU får en vederlagsfri bruksrett til resultatene av oppgaven, inkludert vedlegg til denne, og kan benytte dette til undervisnings- og forskningsformål med de eventuelle begrensninger som fremgår i punkt 8.

8. Utsatt offentliggjøring

Hovedregelen er at studentoppgaver skal være offentlige.

Sett kryss

| | |
|-------------------------------------|------------------------------|
| <input checked="" type="checkbox"/> | Oppgaven skal være offentlig |
|-------------------------------------|------------------------------|

I særlige tilfeller kan partene bli enige om at hele eller deler av oppgaven skal være undergitt utsatt offentliggjøring i maksimalt tre år. Hvis oppgaven unntas fra offentliggjøring, vil den kun være tilgjengelig for student, ekstern virksomhet og veileder i denne perioden. Sensurkomiteen vil ha tilgang til oppgaven i forbindelse med sensur. Student, veileder og sensorer har taushetsplikt om innhold som er unntatt offentliggjøring.

Oppgaven skal være underlagt utsatt offentliggjøring i (sett kryss hvis dette er aktuelt):

| Sett kryss | Sett dato |
|--------------------------|-----------|
| <input type="checkbox"/> | ett år |
| <input type="checkbox"/> | to år |
| <input type="checkbox"/> | tre år |

Behovet for utsatt offentliggjøring er begrunnet ut fra følgende:

Dersom partene, etter at oppgaven er ferdig, blir enig om at det ikke er behov for utsatt offentliggjøring, kan dette endres. I så fall skal dette avtales skriftlig.

Vedlegg til oppgaven kan unntas ut over tre år etter forespørsel fra ekstern virksomhet. NTNU (ved instituttet) og student skal godta dette hvis den eksterne har saklig grunn for å be om at et eller flere vedlegg unntas. Ekstern virksomhet må sende forespørsel før oppgaven leveres.

De delene av oppgaven som ikke er undergitt utsatt offentliggjøring, kan publiseres i NTNUs institusjonelle arkiv, jf. punkt 4, siste avsnitt. Selv om oppgaven er undergitt utsatt offentliggjøring, skal ekstern virksomhet legge til rette for at studenten kan benytte hele eller deler av oppgaven i forbindelse med jobbsøknader samt videreføring i et master- eller doktorgradsarbeid.

9. Generelt

Denne avtalen skal ha gyldighet foran andre avtaler som er eller blir opprettet mellom to av partene som er nevnt ovenfor. Dersom student og ekstern virksomhet skal inngå avtale om konfidensialitet om det som studenten får kjennskap til i eller gjennom den eksterne virksomheten, kan NTNUs standardmal for konfidensialitetsavtale benyttes.

Den eksterne sin egen konfidensialitetsavtale, eventuell konfidensialitetsavtale den eksterne har inngått i samarbeidprosjekter, kan også brukes forutsatt at den ikke inneholder punkter i motstrid med denne avtalen (om rettigheter, offentliggjøring mm). Dersom det likevel viser seg at det er motstrid, skal NTNUs standardavtale om utføring av studentoppgave gå foran. Eventuell avtale om konfidensialitet skal vedlegges denne avtalen.

Eventuell uenighet som følge av denne avtalen skal søkes løst ved forhandlinger. Hvis dette ikke fører frem, er partene enige om at tvisten avgjøres ved voldgift i henhold til norsk lov. Tvisten avgjøres av sorenskriveren ved Sør-Trøndelag tingrett eller den han/hun oppnevner.

Denne avtale er signert i fire eksemplarer hvor partene skal ha hvert sitt eksemplar. Avtalen er gyldig når den er underskrevet av NTNU v/instituttleder.

Signaturer:

| | |
|------------------------------|-----------------------------|
| Studieprogramleder ved NTNU: | |
| Dato: | <i>Jodi Alte</i> |
| Veileder ved NTNU: | <i>Tor/Var Melling</i> |
| Dato: | |
| Ekstern virksomhet: | |
| Dato: 18/01-2024 | <i>Th. Larsen</i> |
| Student: | |
| | <i>Amund R. Steinhilber</i> |
| Dato: 18.01.2024 | |
| Student: | |
| | <i>Lea Arvola Utstein</i> |
| Dato: 18.01.2024 | |
| Student: | |
| | <i>Dina H. Steinstog</i> |
| Dato: 18.01.2024 | |

Appendix C

Project Plan

Prosjektplan del 1 – Gruppe 5

1. Diskuter og gjør en avklaring på problemstilling og eventuelle forskningsspørsmål.

1.1 Background

The use of cloud resources in businesses today is widespread. In Norway, 71% of private enterprises (excluding financial activities) uses some form of cloud computing. This widespread use of cloud resources has more than doubled since 2014 and has potential to grow even more in the future [1]. To manage access to these resources, many businesses have had to generate more and more identity related resources.

An international report from 2022 on *identity sprawl*, the rapid growth of identity-related information a business needs to work with, found 67% of businesses to be experiencing identity sprawl. Businesses generally found identity and access management (IAM) to be of vital importance, but a majority of 61% said that the management process to be too expensive, where 66% cited technical dept as the reason. This culminates in the report findings where 84% of businesses reports to have experienced an identity-related breach, where 67% had one in 2021 [2]. It is likely to believe that many of these breaches could have been prevented if the businesses had been able to focus on IAM and not having to cite technical dept and cost as roadblocks. The FBI reports that something as small as a single business email compromise (BEC) costs a business an average of \$130'000 [3], meaning that businesses which suffer identity related breaches might have to pay large sums of money to clean-up, in addition to any marks which might be left on their reputation.

To further build on the need for secure IAM, the Norwegian Police Security Service (PST) presented in their risk-report for 2023 that they suspected a rise in attacks targeting people. The reason for targeting people is to get to their user identities [4]. This further shows the importance of properly securing IAM within the still growing cloud environments of Norwegian businesses.

Microsoft's cloud platform Azure provides customers with a multitude of options for IAM within the cloud. Through Entra ID for the IAM itself, and Entra ID Protection for management of access and security. Microsoft also provides a SIEM solution in Azure with Microsoft Sentinel, which can be used to manage security incidents as they show up. Sentinel also has automation capabilities which enables businesses to configure automated response to incidents and have queries which look through multiple incidents and tries to look for any correlations between them.

In 2021 Microsoft launched *Content Hub*, a marketplace where organisations could sell rulesets and automations which they had made [5]. Here Microsoft also presents some of their own rulesets which makes it easier for businesses to quickly setup security solution for Entra ID, Azure's IAM. These rules are free for costumers to download should they wish.

1.2 Problem statement

In this thesis we will investigate the rulesets provided by Microsoft in Content Hub for the IAM solution Entra ID and the connected security suite Entra ID Protection. As most businesses reports that technical

dept is an issue for proper IAM, we are choosing to approach the security setup by following Microsoft's recommendations and provided ready-to-go solutions to simulate how an organisation without high-skilled employees could construct a security setup for Entra ID in Azure. This means that we will focus on the enabling of pre-made rulesets in Content Hub for Entra ID and Entra ID Protection and best-practice setup of Conditional Access policies within Entra ID Protection.

We will also attempt to better the rulesets, if needed, by looking at the type of incidents are raised to manual investigation, number of false positives and the automation rules provided.

To test the setup of Sentinel and Entra ID Protection we are choosing a practical approach where we will attempt to simulate actual identity breaches through manually triggering risk detections from Entra ID Protection. As Microsoft claims that only four of their twelve risk detections can be manually triggered, we will investigate the possibility of triggering a wider number of detections.

1.3 Research questions

The main research question for our thesis is the following:

“How well does the rulesets provided by Microsoft in Sentinel Content Hub for Entra ID and Entra ID Protection with a best-practice setup of Conditional Access policies secure an organisation against identity-based threats?”

This research question could be further broken down into four specific questions:

1. Does Sentinel, configured with rulesets for Entra ID and Entra ID Protection provided by Microsoft in Content Hub, provide any additional security features which are not available through the Entra ID Protection dashboard?
2. What alterations needs to be done to the rulesets to minimise false positives and lower the number of incidents marked for manual investigation?
3. How does the use of best-practice Conditional Access policies contribute to the incident handling while using Microsoft's rulesets from Content Hub?
4. Are we able to trigger a risk detection which has not been manually triggered before?

1.4 Effect goals

With these research sub-questions, we hope to achieve the following effect goals:

1. Provide a guide to which alterations are necessary if a Norwegian business is implementing Microsoft's rules for Entra ID and Entra ID Protection in Sentinel, as presented in Content hub.
2. Provide a guide for how to trigger risk detections within Entra ID Protection for businesses to be able to test their own identity security setups within Azure.

1.5 Project goals

When finishing the project, we want to have achieved the following:

1. Have tested the rules provided in Content Hub by Microsoft for Entra ID and Entra ID Protection.
2. Discovered what limits or difficulties which are present when attempting to manually trigger risk detections.
3. Collected a guide for best-practice setup for Conditional Access policies in Entra ID Protection according to current guidelines.

4. Have tested what differences the use of best-practice Conditional Access policies has on the incident detection and handling in the Sentinel setup.
5. Found what more a business will be able to see and do when using Sentinel compared to using the Entra ID Protection dashboard.

2. Søk opp minst to referanser og gi et sammendrag av disse om hvorfor disse er relevante for oppgaven

Our thesis concerns Microsoft products, and how these should be effectively used in our environment to reduce *problems/events* related to *identity*. To gain good knowledge about Microsoft's products, how they are used, and what they are used for, Microsoft's own documentation will serve as a good source for our thesis. Here we can obtain direct information on the two main services we are going to use: Microsoft Entra ID and Microsoft Sentinel (and more, perhaps).

Another aspect our thesis focuses on is *understanding* the type of technology we utilize. To describe the general types of technology, we have found IBM to be a good and relevant source. IBM, like Microsoft, is a reliable actor that has delivered various technologies for decades, giving them a deep understanding of the field. For instance, IBM can explain what a SIEM solution is, and we will gain an understanding of SIEM tools in general, and not just specifically about Microsoft's SIEM solution, Sentinel.

3. Beskriv hvordan dere vil innhente data/resultater”

In order to answer the research questions and address the problem statement, we will utilize quantitative methods for data collection and analysis.

We are going to collect data from Microsoft Sentinel and Content Hub, including log files, alerts, and reports generated by Entra ID Protection, as well as other relevant data sources. In order to do this we will among other things set up our own Sentinel environment. This will provide us with a comprehensive picture of the security situation in the Azure environment and help identify any weaknesses or potential attack vectors.

To evaluate Microsoft's security measures and our implementations, we are going to develop and conduct simulated and real threat scenarios. This will include testing known attack methods (which can be sourced from MITRE) and attempting to trigger new or unexplored risk detections. By exposing the system to various types of threats, we will be able to assess the effectiveness of the security measures and identify any areas for improvement.

In other words, in order to collect data and results we are going to perform different tests and simulations on our system. For instance, A/B testing, also known as split testing, is one way to test that offers an illustrative approach to this. This method involves comparing two or more variations of a system [6], such as one integrated with Sentinel and another without, to gauge their comparative effectiveness.

4. Innhent godkjenning dersom oppgaven krever lagring av persondata

Denne oppgaven trenger ikke innhenting av, eller lagring av, persondata.

5. Definer rapportstrukturen og opprett dokumentet i deres foretrukne teksteditor.

Our thesis is structured into six main chapters: Introduction, Theory, Methods, Results, Discussion, and Conclusion. Each chapter serves a specific purpose in presenting our research findings and analysis.

To organize and structure the thesis, we will use the LaTeX editor Overleaf. The structure is as follows:

- **Abstract:** A brief overview of the purpose of the thesis, the methods used, the main findings, and the conclusions.
- **Sammendrag:** Same content as in Abstract, but written in Norwegian.
- **Contents:** An overview of the thesis's structure and content to help the reader in navigating through the document.
- **Figures:** A list of all figures, tables, and diagrams included in the thesis, along with corresponding numbering and page references.
- **Glossary:** A list of key terms with definitions to assist the reader in understanding specialized terminology used in the thesis.
- **Abbreviations:** A list of abbreviations and acronyms used in the thesis, along with their full meanings.
- **Preface:** A brief introduction describing the background of the thesis, the motivation for carrying it out, and gratitude to any contributors or supporters.

Then the six main chapters:

- **Chapter 1 - Introduction:** In Chapter 1, we introduce the purpose of our thesis and articulate the objectives we aim to achieve. We will also provide some background for the thesis.
- **Chapter 2 - Theory:** Chapter 2 delves into the theoretical framework, exploring concepts such as Microsoft Sentinel and Entra ID to provide a foundational understanding of how they work.
- **Chapter 3 - Method:** Chapter 3 outlines the methodology employed to address our research questions, detailing the data collection, analysis techniques, and evaluation methods utilized.
- **Chapter 4 - Results:** Chapter 4 presents the results obtained from our testing procedures, including any identified findings, trends, or patterns.
- **Chapter 5 - Discussion:** Chapter 5 engages in a comprehensive discussion, contextualizing our results within the broader scope of the thesis's objectives, theoretical underpinnings, and previous research. This chapter also includes reflections on the strengths and limitations of our methodology.

- **Chapter 6 - Conclusion:** Finally, in Chapter 6, we summarize our findings and draw conclusions based on our research questions, offering insights and potential directions for future work and investigation.

Then at the end of the thesis:

- **Bibliography (Bibliografi):** A list of all sources and references used in the report, formatted according to the relevant referencing style (e.g., APA, MLA, Harvard)
- **Appendix:** Supplementary material containing extra information, not essential to the main text but still relevant to the thesis.

Sources:

1. *10966: Buy cloud computing services (per cent), by industry (SIC2007), contents, year and employed.* <https://www.ssb.no/en/statbank/table/10966/tableViewLayout1/>
2. *New Study Reveals Identity Sprawl Plagues Organizations.* 2022. Time of download (26.02.2024). <https://www.radiantlogic.com/news/new-study-reveals-identity-sprawl-plagues-organizations-with-60-percent-reporting-over-21-disparate-identities-per-user/>
3. *Business E-Mail Compromise.* 2015. Time of download (26.02.2024). <https://www.fbi.gov/news/stories/business-e-mail-compromise>
4. *Norsk trusselvurdering.* 2023. <https://www.pst.no/alle-artikler/trusselvurderinger/ntv-2023/>
5. *Introducing Microsoft Sentinel Content hub!.* 2021. Time of download (26.02.2024). <https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/introducing-microsoft-sentinel-content-hub/ba-p/2928102>
6. *A Refresher on A/B Testing.* 2017. <https://hbr.org/2017/06/a-refresher-on-ab-testing>

Prosjektplan del 2 – Gruppe 5

6. Hvilke ressurser må dere ha tilgang til for å kunne gjennomføre oppgaven? (lisenser, utstyr...)

To answer the thesis, we need access to a tenant in Microsoft Azure, so that we get access to Microsoft Entra ID and Microsoft Sentinel. We have already been handed over a subscription from NTNU with help from our supervisor. The license we have is the Microsoft 365 E5 Developer, and in addition we also have a subscription which is an Enterprise Agreement (agreement between Microsoft and NTNU). By having the subscription and license we have we get access to both Microsoft Entra ID, Microsoft Sentinel, and more if needed.

7. Hvilke sentrale aktiviteter har dere i prosjektet?

I prosjektet har vi sentrale oppgavene av å skrive teori, metode og diskusjon, og innhenting av resultater. Innenfor hver av disse har vi skapt en noe detaljert liste over arbeidsoppgaver vi ser for oss må bli gjort. Dette har vi satt opp i et Gantt diagram her:



Her er en mer detaljert forklaring til hva som inngår i hver av de forskjellige hovedoppgavene (underveis kan det være at vi identifiserer nye oppgaver og legger til, eller fjerner overflødige oppgaver):

- Sette opp miljø
 - Sette opp Sentinel på Tenant
 - Sette opp brukere
- Skrive teori
 - Teori: Sentinel

- Teori: Content hub
 - Sette opp regler i Sentinel / Content hub
 - Manuell inspeksjon av content hub regler og forklar hvordan hver av dem fungerer.
- Teori: Entra ID
- Teori: Conditional Access
- Teori: Risk Detections
- Teori: Identity and Access Management (IAM)
- Teori: Entra ID Protection
- Funnet ut hvilke ekstra bonuser en bedrift generelt vil få ved å bruke Sentinel kontra Entra Identity Protection
- Skrive metode:
 - Skrive metode for triggering av risk detections
 - Simulere bruksmønster (logge inn på brukere)
 - Teste triggering av risk detections (bare slik at vi vet at det er mulig)
 - Sette opp CA i ME-ID og skrive begrunnelse for hvorfor dette er best-practice
 -
- Innsamling av resultater:
 - A/B test: med/uten Sentinel regler fra content hub
 - A/B test: med/uten CA
 - Finne hvilke problemer som oppstår ved manuell triggering av risk detections
- Skrive diskusjon:
 - Skrive en guide til forbedringer som må bli gjort på regelsettene fra content hub
 - Skrive en guide til best-practice setup for CA i henhold til guidelines
 - Skrive en guide til hvordan man kan gjennomføre testing
- Forbedring av oppgaven
 - Kvalitetsikre oppgaven

Vi har fordelt disse oppgavene inn i "arbeidssiloer" hvor hver silo er uavhengig av hverandre. Dette gjelder frem til "innsamling av resulater" hvor det er stort overlapp og det er påkrevd at vi har gjort ferdig betraktelige deler av metode og teori.

8. Hvordan fordeles ansvaret mellom gruppemedlemmer?

Vi fordeler arbeidet mellom oss slik at én person tar for seg silo 1 (Sentinel og Content Hub), en annen tar for seg silo 2 (Entra ID) & 3 (Entra ID Protection u/ risk detections), og den tredje tar for seg silo 4 (Entra ID Protection, med fokus på risk detections).

Silo og ansvarlig person (slik som vi har fordelt nå):

1. Dina
2. Lea
3. Lea
4. Amund

9. Hvilke milepæler har prosjektet? (når skal ulike deler være klare)?

Som man kan se i Gantt diagrammet over har vi identifisert følgende milepæler for prosjektet:


1. Milestone: Regler fra Content Hub implementert (uke 12)
2. Milestone: Prosedyre for test av risk detections klar (uke 14)
3. Milestone: Best-practice CA implementert (uke 14)
4. Milestone: Alle brukere ferdig med "eksistensperiode" (uke 15)
5. Milestone: Ferdig førsteutkast (uke 19)
6. Milestone: Levere oppgave (uke 20)

Alle disse markerer store punkter i oppgaven vår som vi kan måle vår fremgang ut ifra. Vi forventer at disse er ferdige til den uken de slutter mot (står i parentes i listen over).

Milepæl nummer 4 påpeker at perioden, for hvor lenge en bruker må ha eksistert for å kunne trigge noen risk detections i Entra ID Protection, har gått og at alle brukere kan bli testet på, gitt at de ble skapt innen fristen for *setup av miljø*.

Appendix D

Gantt Chart








Appendix E

Time Table

Sluttrapport total timetall summert pr deltaker pr aktivitet

| Aktivitet | Dina | Lea | Amund | Total gruppearbeid sum pr aktivitet |
|-----------------------------------------------|--------------|--------------|------------|-------------------------------------|
| Gruppemøte | 49.5 | 48 | 48 | 145.5 |
| Informasjonsinnhenting (forprosjekt) | 42.5 | 35.5 | 26 | 104 |
| Skrive metode | 31 | 18.5 | 40.5 | 90 |
| Skrive teori | 22.5 | 46.5 | 49 | 118 |
| Innsamling(testing) og skriving av resultater | 57 | 70.5 | 17 | 144.5 |
| Administrativt | 13.5 | 14.5 | 21 | 49 |
| Skrive diskusjon | 37 | 18 | 15 | 70 |
| Forbedring av oppgave | 62.5 | 57 | 61 | 180.5 |
| Setup av miljø | 24 | 3 | 0.5 | 27.5 |
| Total arbeid pr student | 339.5 | 311.5 | 278 | 929 |

| | |
|--------------|------------|
| Startdato | 08.01.2024 |
| Slutt dato | 21.05.2024 |
| Antall dager | 134 |



Appendix F

Minutes of Meetings with Supervisors

Møteinnkalling til oppstartsmøte 10.01.2024

Dato og tid: Torsdag 18.01.24 kl. 10:00 – 12.00

Sted: Sluppenvegen 17A rom 2A

Følgende personer innkalles:

Dina

Lea

Amund

Tor Ivar (veileder)

Thor (ekstern)

Tobias (ekstern)

Cate (ekstern)

Peder (ekstern)

Agenda:

Sak nr 01 Gjennomgå oppgaveinformasjonen vi har til nå

Sak nr 02 Gjennomgå standardavtale

Sak nr 03 Praktiske spørsmål til veileder og eksternt team

Møtet planlegges avsluttet ca kl. 12.00

Ta kontakt med undertegnede dersom du ikke har anledning til å komme

Mvh

Dina

Referat fra oppstartsmøte 18.01.2024

Dato og tid: 18.01.24 kl 10:00-12:00

Sted: Sluppenvegen 17A, rom 2A

Til stede: Dina, Lea, Amund, Tor Ivar (veileder), Peder (ekstern), Cate (ekstern), Tobias (ekstern) og Thor (ekstern)

Frafall: Ingen

Ordstyrer: Amund

Sak nr. 1 – Gjennomgå oppgaveinformasjonen vi har til nå

Vi oppsummerte oppgaven, og kom med innspill på hvordan vi har tenkt å løse den

- Hvordan bruke Sentinel og det den er bra til
 - Bruk Sentinel til sikkerhetsmekanismer
- Sammenligning med regelsettene Microsoft utgir i Content Hub, om de er effektive, og egentlig bra nok
- Hva skal vi sette opp som mål, for å vurdere om reglene er bra nok
 - Hva er disse reglene basert på, hvordan har vi kommet frem til kriteriene
 - Samle inn data for å lage kriterier basert på håndfast informasjon

- MS Entra ID
 - Sette opp aksessmekanismer
 - Bruk til identitet
- Eventuelt sette nye regler for å tette hull, eller fjerne regler som ikke er nødvendig
- Implementeringen av Zero-trust i oppgaven, zero-trust er en måte å designe infrastrukturen på.

- Kubernetes (AKS)
 - Skal vi ha med dette? Kan dette introdusere en helt ny oppgave?

- Finne en testmekanisme som er god
 - Finnes det et slikt verktøy?

- Tester:
 - Hva skal testes?
 - Se på hva som faktisk brukes i markedet. Det som blir brukt mest er det som burde testes. Dette blir en del av datainnsamlingen.
 - Dette kan være styrende for hva vi ønsker å teste, ettersom dette er hva markedet bruker mest.
- Viktig å begrunne hva man tester og hvorfor.

- Hvilken sektor:
 - Dette må man velge mens man gjør innsnevring nå i starten.
 - Vi må bare lese litt her og der og se litt hvordan hver av sektorene er. Ut ifra hvordan vi ønsker å utforme oppgaven så kan vi bestemme en sektor ut ifra dette.

- Viktig å argumentere for hvorfor.
 - Vi har sett på dette, så på dette og derfor valgte vi dette.
- Se på MS cloud security benchmark for å se på om tenant er “compliant” med kriterier for de forskjellige sektorene.
 - Sjekk ut parview (purview? paraview?)

Sak nr. 2 - Gjennomgå standardavtale

Gjennomgikk standardavtalen med veileder og eksternt team (Tietoevry)

- Lagt i teams for signering, gjort til to do

Sak nr. 3 – Praktiske spørsmål til veileder og eksternt team

Spørsmål til veileder og eksternt team (Tietoevry):

Vi hadde en del praktiske spørsmål som vi trengte et svar på, vi kom frem til;

- Hvordan skaffe tilgang til Azure?
 - Sett opp devportal på tenant, inviter inn veileder, Tor Ivar, før han kan gi oss tilganger til det vi trenger
- Hvordan dele timelisten med veileder
 - Lagt til som fil i Teams, her kan Tor Ivar sjekke filen når han enn måtte ønske!

Møteinnkalling - Bachelorgruppe 5

Dato og tid: Torsdag 01.02.2024 kl. 10:00 – 11.00

Sted: Adolf Øien-bygget rom G316

<https://use.mazemap.com/#v=1&config=ntnu&zlevel=3¢er=10.398855,63.415191&zoom=18&sharepoi=182244&campusid=1>

(eventuelt digitalt over Teams)

Følgende personer innkalles:

Dina

Lea

Amund

Tor Ivar (veileder)

Agenda:

Sak nr 01 Godkjenning av referat fra oppstartsmøte

Sak nr 02 Gjennomgå av hva vi har gjort til nå

Sak nr 03 Praktiske spørsmål til veileder

Møtet planlegges avsluttet ca kl. 11.00

Ta kontakt med undertegnede dersom du ikke har anledning til å komme

Mvh

Dina

Møtereferat 01.02.2024 - Bachelorgruppe 5

Dato og tid: Torsdag 01.02.2024 kl. 10:00 – 11.00

Sted: Adolf Øien-bygget rom G316

<https://use.mazemap.com/#v=1&config=ntnu&zlevel=3¢er=10.398855,63.415191&zoom=18&sharepoi=poi&sharepoi=182244&campusid=1>

(eventuelt digitalt over Teams)

Følgende personer innkalles:

Dina

Lea

Amund

Tor Ivar (veileder) - over teams

Agenda:

Sak nr 01 Godkjenning av referat fra oppstartsmøte

Tor Ivar (veileder) godkjente referat av oppstartsmøte

Sak nr 02 Gjennomgå av hva vi har gjort til nå

- Snakket om funn fra sektor
- Hørt angående intervju, use case og gjennomføring av brukertest (eventuelt senere i oppgaven)

Sak nr 03 Praktiske spørsmål til veileder

- Hvor faglig begrunnelse kreves?
 - Gjerne litt data/statistikk, men også forhåndssatte krav fra oppgavegiver, tilgjengelighet fra skolen o.l.
- Hvordan sitere norske sitater i engelsk oppgave?
 - Usikker, laget to do til Tor Ivar (veileder), slik han kan få hørt med resten av fagstab.
- Hvordan kildehenviser, hvor mye skal referes til?
 - Gjerne kilde på alt, selv selvopplagte påstander
- Hvilken subscription får vi i ME-ID?
 - Får P2 og E5 gjennom skolen, Tor Ivar må legges til for å gi subscription i vår eksisterende profil
- Når passer møter med Tor Ivar (veileder)
 - Torsdager klokken 10 fungerer!

Møtet avsluttet kl. 11.00

Møteinnkalling - Bachelorgruppe 5

Dato og tid: Torsdag 15.02.2024 kl. 10:00 – 11.00

Sted: Adolf Øien-bygget rom G316

<https://use.mazemap.com/#v=1&config=ntnu&zlevel=3¢er=10.398855,63.415191&zoom=18&sharepoi=182244&campusid=1>

(eventuelt digitalt over Teams)

Følgende personer innkalles:

Dina

Lea

Amund

Tor Ivar (veileder)

Agenda:

Sak nr 01 Gå gjennom spørsmålene angående sitering

Møtet planlegges avsluttet ca kl. 11.00

Ta kontakt med undertegnede dersom du ikke har anledning til å komme

Mvh

Dina

Deltagere: Dina, Lea, Amund, Thor.

Sak nr. 1:

Thor: Dere kan hente dere lånekort nå. Dere skal nok få beholde dem over lengre tid ettersom dere har blitt registrert i workplace.

Sak nr 2:

Thor: Har dere tenkt på å ha intervju med eksperter på tema for å få et større perspektiv? Dette kan være bra for å få et kult og nytt perspektiv for oppgaven. Kornelius&Co har gjort dette allerede.

- Dette har vi diskutert med Tor Ivar og det har vist seg å være litt ekstra administrativt arbeid. En stor del av arbeidet handler om riktig metode så vi må også unngå å ikke.
- En del av selve testingen
- Det kan være mulig å se på hva slags alternative møter vi kan sette opp som ikke er intervjuer.
- Spørsmål til fremtiden: Kan det være mulig å skaffe informasjon uten å gjøre det som et ordentlig intervju?

Møte avsluttet 10:30.

Møteinnkalling - Bachelorgruppe 5

Dato og tid: Onsdag 21.02.2024 kl. 09:00 – 10.00

Sted: Digitalt over Teams

Følgende personer innkalles:

Dina

Lea

Amund

Tor Ivar (veileder)

Kaller inn til teams-møte onsdag 21.02 klokken 09:00! Møte vil bli avholdt over Teams

Agenda:

- Hvordan prate med ekspertise i Tietoevry og bruke dette i vår rapport – intervju, eller noe mindre formelt?
- Lande på en problemstilling – tips og innfall på problemstillinger (og research questions) vi allerede har definert
- Eventuelt andre spørsmål som har dukket opp

Møtet planlegges avsluttet ca kl. 10.00

Ta kontakt med undertegnede dersom du ikke har anledning til å komme

Mvh

Dina

Møtereferat 21.02.2024 - Bachelorgruppe 5

Dato og tid: Onsdag 21.02.2024 kl. 09.00 – 10.00

Sted: Digitalt over Teams

Deltagere: Dina, Amund, Lea og Tor Ivar

Agenda:

Sak nr. 1:

Hvordan prate med ekspertise i Tietoevry og bruke dette i vår rapport – intervju, eller noe mindre formelt?

Den andre gruppen tok opptak med transkribering(?) slik at de kunne sitere til transkriberingen. Veileder Tor Ivar sier at det kan være lurt å finne en annen kilde som bekrefter det kilden sier, slik at vi kan henviser til flere kilder som bekrefter det samme.

Sak nr. 2:

Lande på en problemstilling – tips og innfall på problemstillinger (og research questions) vi allerede har definert

Finner forslag under Notater -> Mulige problemstillinger

Tor Ivar tanker: Problemstilling kan være like lang som *ett* research question, så kan research questions heller bli delt inn i *tre*. Det var veldig mye tekst, noe kan man spille mye på, andre ville han utelukket. Noen research questions sliter han også med å forstå hvordan man skulle ha besvart, eller hvordan man skulle ha funnet informasjon for å besvare spørsmålet.

Oppsummert:

En mer innholdsrik problemstilling (som forklarer hensikten med oppgaven litt bedre), som består av tre forskningsspørsmål vi ønsker å besvare

Mer kjøtt i problemstillingen, mindre ull i forskningsspørsmålene

Sak nr. 3:

Eventuelt andre spørsmål som har dukket opp

- 1) Hvordan skrive sitat inn i oppgaven hvis de er skrevet på norsk?

Veileder har pratet med Joakim (emneveileder?), noter/poengter at man har oversatt kilden, skriv på engelsk, og henvis til kilden.

Møteinnkalling - Bachelorgruppe 5

Dato og tid: Tirsdag 27.02.2024 kl. 10:00 – 11.00

Sted: Digitalt over Teams

Følgende personer innkalles:

Dina

Lea

Amund

Tor Ivar (veileder)

Agenda:

- Problemstilling
- Andre spørsmål

Mvh,
Lea

Møtereferat 27.02.2024 - Bachelorgruppe 5

Dato og tid: Tirsdag 27.02.2024 kl. 10.00 – 11.00

Sted: Digitalt over Teams

Deltagere: Dina, Amund, Lea og Tor Ivar

Agenda:

Sak nr. 1:

Gjennomgang av ferdigstilt problemstilling og endringer som er gjort.

- Tor Ivar la igjen kommentarer på problemstillingen før møtet
- Endringene som er gjort ser bra ut
- Forskningsspørsmål 2 ble endret på og vi fant en bedre formulering på spørsmålet.

Videre gikk vi gjennom effektmålene og resultatmålene. Der kom vi fram til at vi burde endre på effektmålene slik at de er målbare, og ikke bare en guide.

Sak nr. 2:

Gjennomgang av diverse spørsmål som vi hadde

- Hvor detaljert skal gantt diagrammet være?
 - Trenger ikke være veldig detaljert, gjør det veldig enkelt
- Må vi dokumentere setup?
 - Vi kan forklare hva vårt oppsett er, men trenger ikke og forklare steg for steg. Skriv heller en kortfattet beskrivelse oppsummert av hva vi har gjort. Trenger ikke å gå i detaljer. Nevn hva det er å hva det gjør. Selve oppsette er under metode-delen der vi forklarer hva oppsette er. I tillegg kan vi ta utgangspunkt i at leseren har grunnleggende forståelse for hva Microsoft cloud er.
- Burde teorien skrives først?

- Teorien burde være ferdigskrevet siden den skal kunne knyttes opp mot det seinere i oppgaven. Skriv mest mulig ferdig, men vær obs på at man kanskje må gå tilbake å skrive litt ekstra.

Møteinnkalling - Bachelorgruppe 5

Dato og tid: Torsdag 7.03.2024 kl. 10:00 – 11.00

Sted: Sluppen 17A rom 2A

Følgende personer innkalles:

Dina

Lea

Amund

Thor (ekstern)

Tobias (ekstern)

Cate (ekstern)

Peder (ekstern)

Nå er forprosjektperioden over og vi ønsker å komme med en oppdatering på hva problemstillingen vi har landet på går ut på og hvordan veien fremover ser ut.

Agenda:

- Presentasjon av problemstilling
- Veien fremover
- Eventuelle spørsmål eller tanker om problemstilling og oppgave

Møtereferat 7.03.2024 - Bachelorgruppe 5

Dato og tid: 7.03.24 kl 10:00-11:00

Sted: Sluppenvegen 17A, rom 2A

Deltagere: Dina, Lea, Amund, Peder (ekstern), Cate (ekstern), Tobias (ekstern) og Thor (ekstern)

Sak nr. 1 – Presentasjon av oppgave

Vi oppsummerte oppgaven og hva vi har kommet fram til i en presentasjon. Her gikk vi gjennom problemstillingen, forskningsspørsmålene, og prosjektplanen fremover.

Kommentarer om Sentinel-delen av oppgaven:

- Anbefaler oss at vi må vise nytten av å bruke Sentinel contra ikke (det finnes løsninger i dag som ikke benytter Sentinel som fungerer helt fint). Vi burde vise til fordelene ved å bruke Sentinel.

Kommentarer om CA- og ID Protection-delen av oppgaven:

- I dag er det “best practice” å bruke CA i stedet for ID Protection. Det er ikke noe som konfigureres i ID Protection, dette gjøres i CA. Før var det greit å bruke ID Protection, men det er ikke det som anbefales lengre eller sees på som “best practice”.
- Anbefalte oss å bruke CA som Microsoft kommer med, slik at vi ikke må konfigurere noe selv. (Men husk å begrunne **hvorfor** dette er best practice”)
- Hva er egentlig “best practice”? Vi må definere hva det er for oss i oppgaven, er dette Microsoft best practice? Eller er det community best practice?

Sak nr. 2 - Spørsmål

Gjennomgang av diverse spørsmål/kommentarer vi og de hadde:

- Til fs 2: Vi burde skrive om siste delen til forskningsspørsmål (minimer til falske uten true positive?) da dette kan være viktig å få med denne typen drøfting når vi skal drøfte.
- Til fs 1: Ville det hørtes mer teoretisk hvis vi hadde endret “additional” til “other”?
- Hva menes med “Security features”? Hvilke er det vi tenker å sammenligne? Vurder å endre det til tools.
- De 3 først forskningsspørsmålene er greie nok i seg selv, så ha nr. 4 med som en bonus.
- Møter fremover:
 - Vi kan bestemme og innkalle når vi har behov. Vi burde også helst spørre de spørsmålene vi har med en gang i stedet for å samle de opp.
- Når det kommer til generelle tips til skrivning o.l. kan dette også tas underveis.
- Metoden er kun teorien på forskningsspørsmålene. Hvordan skal vi finne og løse forskningsspørsmålene. Hvilken data som skal hentes, hvordan den er hentet, og hvordan den har blitt analysert/behandlet osv.

Møteinnkalling - Bachelorgruppe 5

Dato og tid: Tirsdag 14.03.2024 kl. 10:00 – 11.00

Sted: Digitalt over Teams

Følgende personer innkalles:

Dina

Lea

Amund

Tor Ivar (veileder)

Ny innkalling her ettersom det ble litt kluss med forrige.

Agenda:

- Få tilbakemelding på:
 - Kapittel 1: Introduksjon
 - (deler av) kapittel 2: Teori
- Gå igjennom spørsmål

Mvh,
Amund

Møtereferat 14.03.2024 - Bachelorgruppe 5

Dato og tid: 14.03.24 kl 10:00-11:00

Sted: Teams

Deltagere: Dina, Lea, Amund, Tor Ivar (veileder)

Sak nr. 1 – Tilbakemelding på skriving

- Kapittel 1:
 - Tilbakemeldinger på det vi har skrevet er lagt til i kommentar i tasken vi ga til veileder.
 - Vi burde omformulere scope og limitations litt her og der (til nå virker det litt hakkete og oppdelt). Veileder skjønner hva som står der, men får følelsen av at vi skyver unne en del ting (dette er første inntrykk når han leste gjennom første gang). Det er forståelig hvorfor det er skrevet, men det blir veldig mye som står ute av scope.
 -
- Kapittel 2 - Teori:
 - Vi trenger ikke nødvendigvis å ha en introduksjon til hver seksjon i teori delen av oppgaven, men dette er noe vi selv kan bestemme.

Sak nr. 2 - Spørsmål

Gjennomgang av diverse spørsmål/kommentarer vi og de hadde:

- Hvor kan f.eks. forklaring av brute force puttes?
 - Her kan vi anta at leseren vet hva dette er, så vi trenger kun å skrive kort om hva det er der vi skriver om trusselbildet i dag eller om vi har en egen teori del til dette.
- Vi kan fortsette å sende denne mengden tekst til veileder slik at kan lese gjennom og gi tilbakemeldinger.

- Vi trenger ikke å sette opp et møte for hver gang veileder har lest gjennom tekst. Skriftlig tilbakemeldinger fungerer helt greit.
- Hvis det er noe som skal gås gjennom på møte er det greit at vi gir teksten som skal leses gjennom tidlig/god tid i forkant (hvis ikke kan vi ikke forvente at vi har fått tilbakemeldinger)
- Når vi har endret på det vi har fått tilbakemelding på så trenger vi ikke å sende teksten til veileder på nytt.
- Når vi skal gjennomføre tester kan vi fint bruke verktøy som vi finner på nett, vi må bare huske å skissere i oppsettet hva som er brukt for å gjennomføre testene. Vi trenger ikke å forklare i detaljer, bare litt hva det utfører.
- Vi kan bruke AI (grammarly, chatgpt) som språkvask.
 - Så lenge det er dine ord som er skrevet på engelsk, så kan du bruke chatgpt f.eks. til språkvask (customized chatgpt, customized innstillinger).
 - Ha med begrunnelse på hvorfor verbet f.eks. er feil, sånn at man lærer.
 - Ikke bruk ord som du ikke forstår eller er komfortabel å bruke.
- Fotnote, glossary-list:
 - Der det føles hensiktsmessig kan man bruke fotnote, men da må man være konsekvent og bruke det til hva det er til. Dette er noe vi i gruppen kan bestemme.
 - Hvis det er noe som ikke gir mening å forklare i oppgaven, skriv det i fotnoten.
 - Glossary list: gloseliste skal brukes for ord og definisjoner for spesifikke ting relevant til oppgaven
 - Fotnote: ekstra informasjon, som egentlig ikke er relevant til oppgaven. Brukes for å unngå å plutselig bryte opp det vi skriver.
- Hva skal i gloselisten og hva skal i teori delen?
 - Les oppgaven som er lastet opp og se på oppsettet der, den er veldig bra. Ikke tenk på innholdet, men mer på oppsettet. (oppgaven ligger under files på teams)
 - Formatet på oppgaven er et gått eksempel på format vi kan følge.

Møteinnkalling - Bachelorgruppe 5

Dato og tid: Onsdag 10.04.2024 kl. 10:00 – 11.00

Sted: Digitalt over Teams

Følgende personer innkalles:

Dina

Lea

Amund

Tor Ivar (veileder)

Hei! Kaller inn til nytt teams-møte onsdag 10.04 klokken 10:00. Møte vil bli avholdt over Teams.

Agenda:

- Tilbakemelding på:
 - Kapittel 3: Metode
 - (kanskje) kapittel 2: Teori

- Gå gjennom spørsmål

Med vennlig hilsen,

Lea

Møtereferat 10.04.2024 - Bachelorgruppe 5

Dato og tid: 10.04.24 kl 10:00-11:00

Sted: Teams

Deltagere: Dina, Lea, Amund, Tor Ivar (veileder)

Sak nr. 1 – Tilbakemelding på skriving

- Kapittel 3 – Metode:
 - Veldig mye bra. Noen ting:
 - Fin introduksjon med klare rammer.
 - Veldig mye under analytic rules, føles som en evig opprømsing.
 - Så lenge det er relevant, så er det ikke feil
 - Kan la det stå slik det er inntil videre, så kan vi se senere om det burde endre på.
 - Snakker om mfa konfigurasjon
 - Ikke ta med at den er skrudd av fra før. Ikke ha med noe som kan forvirre leser
 - Kapittel 3.4: Litt usikker på om det burde være litt mer bakgrunn for det.
 - Metode handler jo om hvordan har vi samlet inn data.
 - Få med valg og metode: Savner litt mer om styrker og svakheter til metodene vi har brukt. Klarer vi å finne styrker og svakheter med fremgangsmåten og komme med en selvevaluering.
 - Dette kan være med i metode. En egen vurderingsdel i metode kapitlet.
 - Skjønner vi at når vi valgte den fremgangsmåten at det er svakheter/styrker med den? Hva er styrkene ved å gjøre det? Hva er svakhetene ved å gjøre det?
 - Ha med mer helhetlig på slutten under en samlet overskrift. Ikke ha en for hver enkelte.
 - For å vise at vi vet at de valgene vi har tatt påvirker resultatet.
 - Testen er ikke relevant, metoden er at vi viser til hvordan metode vi har brukt for å samle inn data.
 - Helhetlig er det veldig bra.
 - Oppsetts-messig er det veldig bra, fin struktur
- Kapittel 2 - Teori:
 -

Sak nr. 2 - Spørsmål

Gjennomgang av diverse spørsmål/kommentarer vi og de hadde:

- Endret vinkling fra risk detections, til mer generelle angreps metoder.
 - En fornuftig løsning.
- Er det greit å også si at workloads er utenfor oppgaven og bare se på risk detections for brukere? *Da kan det være at vi må skrive om introduksjonen.*
 - Ja, tror også det å fokusere på brukeridentiter kan være interessant.
- Hva av det vi har skrevet i metode er det som skal stå der?
 - Beit litt merke i introen til 3.4, den er fin å ha med, men prøv å bak litt mer rundt den.
- Hvordan skal man referere til github, skal man ha refere til brukernavn og slikt?
 - Finnes et verktøy for å referere til github repository? **Zotero**, kan sendes til bibtex
 - Ang. Verktøy for brute force: kan det være lurt å begrunne hvorfor vi har brukt det? Vær bevist på valg av verktøy.
- Må vi referere til det mest spesifikke stedet hvor man har funnet info eller kan man ta det fra en samleside hvor man samler mye informasjon?
 - Spesifk vs mer generell?: Litt usikker, men vil anta at det er bedre å referere til den mer spesifikke på hver av de.
 - Analytic rules: Må hver enkelt side refereres til? Nei, høres mer fornuftig ut å referere til en samlet side.
 - Se an. Ikke ha 40 ref. Til 40 regler, men er det mer tekstspesifikt og teksten er mer utfyllende, så ta det dit.
- Når man referer til en kilde som er skrevet av mange (ms docs, github osv), referer man til den siste som bidro, den som skrev det første gang, eller til eier av siden?
- Det finnes ingen anbefalinger for playbooks i Sentinel, kan vi bare sette opp det vi tenker er logisk, må vi da begrunne valgene våre med hensyn til teori?
 - Basert på mangel på anbefaling, ha det med som en side-note: her er noe som kan gjøres med hendelsene som en automatisert respons, mens siden vi ikke har noe som er anbefalt så har vi det ikke med. Oppgaven fokuserer mer på innsamling.
 - Hva skal vi svare på? Hvis playbook er med på hva som må gjøre for å kunne besvare problemstilling, men hvis ikke må vi begynne å revurdere.
 - Hvis det skal være med, så må det gjøre noe mer med det.
 - Kan begrunne oppsett av playbooks ved grunn i hvilke mitigations som anbefales fra mitre.
 - Hvis vi skal ha med playbook, så må vi kunne begrunne det og det må være noe bak det.
 - Endre rulesett til solution
 - Fortsetter å se på deteksjon, altså innsamling av data- ikke respons.
 - Ca: hvordan påvirker bruken av ca deteksjon.
- Må vi navngi testbrukerene i metode?
 - Skisser mer testing fremfor brukerne som er brukt.

- Hvordan vi gjennomfører testingen. 3.2: Noen har spesifikke steg, men andre bare referer til teori. Hva er foretrukket?
 - God praksis å referere til teori-kapittelet.
 - Står noe allerede i teori, så er det bare å referere til det.

Møteinnkalling - Bachelorgruppe 5

Dato og tid: Mandag 13.05.2024 kl. 10:00 – 11.00

Sted: Digitalt over Teams

Følgende personer innkalles:

Dina

Lea

Amund

Tor Ivar (veileder)

Inviterer til teams møte 13.05 klokken 10 for tilbakemelding på innsendte biter, pluss eventuelle spørsmål som måtte forekomme!

Mvh

Dina Hagen Steinskog

Møtereferat 13.05.2024 - Bachelorgruppe 5

Dato og tid: 13.05.24 kl 10:00-11:00

Sted: Teams

Deltagere: Dina, Lea, Amund, Tor Ivar (veileder)

Sak nr. 1 – Tilbakemelding på skriving

- I metoden står beskrivelsen av alle analytic rules, skulle denne forklaringen heller ha stått i appendix? (noen av disse beskrivelsene benyttes i diskusjonen)
 - Ja flytt til appendix
- I diskusjonsdelen står det en beskrivelse og sammenligning mellom de to dashbordene (risk detection og sentinel incident). Burde denne stå der de står eller burde de flyttes på/fjernes?
 - Bør det at det var en større forskjell en forventet kanskje drøftes mer?
 - Det med dashbordene er greit å ha med, men diskutere resultatene enda mer.
 - Vi nevner ikke hva vi hadde forventet.
 - Bevisstgjør leseren om at vi har tenkt på det, og viser at vi aksepterer at vi ikke vet og vi ikke har en forståelse for det.

- Så lenge det med maskingjøren er nevnt sånn overordnet eller i en setning her og der, så er det ikke så farlig at det er med.
- Diskusjon:
 - Vurder en kort intro til kapittelet (hva er det vi skal ha med)
- Annet:
 - Surr på formulering i kap. 3.1.2
 - Kap 3.2.: Litt muntlig shwong, gjør litt mer formelt.
 - Unngå å bruke “as described”, bruk heller “as mentioned”
 - 3.5: analytic rules, allerede diskutert – referer til appendix
 - 3.6.1: Bra med evaluasjon!
 - 3.6.2: Skriv om den første setningen.
 - 3.6.3: Første setning litt tung setning.
 - 3.6.4: Bra 😊

Sak nr. 2 - Spørsmål

- Anbefales det å ha en oppsummering av resultatene (på starten/slutten av resultat kapittelet)?
 - Nei, resultatet kan være som det er
- Er det nødvendig å ha med alle tabellene i resultatene, eller blir det litt overflødig (vi påpeker i metoden at det er de tre spørsmålene som er viktige og det vi benytter oss av, tabellene er ikke mye i bruk heller)?
- Presenteres resultatene på en god måte?
 - Tabellen vi presenterer er god og fin, tror det kan være av interesse og ha med.
 - Resultatene er output fra noe vi har gjennomført i metoden, resultatet skal bare presentere resultatet as is; her har du det vi har funnet. Her er det vi fant på den og den testen.
 - Hva er det den her testen har som mål og finne ut?
 - Er de tre spørsmålene er en bra måte å gjøre det på?
 - Presentasjonen av resultatene er ok slik det er satt opp nå, fint med de tre kategoriene
 - De som leser gjennom forventer ikke en forskningsartikkel, teller mer om resultatene er presentert på en fin måte.
 - Bekrefter at resultatene er presentert bra, det gir mening
- Appendix:
 - Skal timeliste, prosjektplan osv. Inn i rapporten sin appendix?
 - Det skal inn i et eget appendix i den endelige innleveringen
 - Del opp i a, b, c osv..
- Presentasjon:
 - Hovedessens, ikke detaljer, overordnet bilde, the big picture
 - 15 minutter
 - Kommer mer info på blackboard



 **NTNU**

Norwegian University of
Science and Technology