UNIVERSITY OF TURKU

Increasing resilience in privileged access management

Cyber Security

Master's Degree Programme in Information and Communication Technology

Department of Computing, Faculty of Technology

Master of Science in Technology Thesis

Author:

Tomas Kopra

Supervisors:

Antti Hakkala (University of Turku)

Petri Sainio (University of Turku)

December 2023

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

Subject: Cyber Security
Programme: Master's Degree Programme in Information and Communication Technology
Author: Tomas Kopra
Title: Increasing resilience in privileged access management
Number of pages: 50 pages, 0 appendix pages
Date: December 2023

This Master of Science in Technology thesis, developed in collaboration with a target company, focuses on increasing resilience and disaster recovery planning for a privileged access management tool. The research was conducted using online sources and supplemented by available frameworks and best practices while working for the target organization.

The thesis explorers several critical questions regarding privileged access rights: their nature, necessity to secure them, appropriate protection mechanisms, and ensuring the resilience of the protection mechanisms during potential disaster recovery scenarios.

The research adopts design science research methodology, commencing with a literature review of identity and access management. The thesis progresses by identifying and assessing possible threat, incident and disaster scenarios for privileged access. The research then presents the most relevant scenarios and solution to enable resilience through high availability. The solutions is then evaluated. The research culminates in a conclusion that answers the set research questions.

Keywords: Disaster Recovery, Privileged Access Management, Identity and Access Management, High Availability.

Table of contents

Acknowledgements

I would like to express my deepest gratitude to all of the members of the target organization's extended IAM team, for providing the means to write this thesis as a part of the team. Many thanks to other supporters and motivators, such as my former teachers and professors, that helped me to find a career in cybersecurity. A special thank you to my father for all the years of work and support that empowers my work.

Glossary

Attack Surface: Attack surface refers to all possible points and methods through which an attacker can exploit vulnerabilities in an IT environment.

Brute-force attack: Brute-force attack is a trial and error method used by attackers to gain unauthorized access to a system. In this type of attack, the attacker will systematically try all possible combinations of usernames and passwords.

Credentials: Credentials typically refer to pieces of information that confirm an identity or authority to access a system.

DNS (Domain Name System): DNS is a hierarchical decentralized naming system for computers, services, or any resource connected to the internet or a private network. It translates user-friendly domain names into IP addresses.

Framework: a security framework refers to a structured set of guidelines, best practices, and standards designed to help organizations manage and improve their overall security posture.

HTTPS (Hypertext Transfer Protocol Secure): HTTPS is an extension of the standard Hyper Text Transfer Protocol (HTTP) used for secure communication over a computer network, particularly the open internet.

Heartbeat: Heartbeat refers to a periodic signal or message exchanged between devices or services to confirm their operational status. Commonly used in high availability setups or clustering environments.

Clustering: Clustering, refers to the use of multiple interconnected computers or servers and distributing the workload among the cluster nodes, which are the individual computers or servers.

SIEM: Security Information and Event Management, a comprehensive cybersecurity service that combines security information management and security event management, where security data from various sources is collected, analysed, and correlated to detect and respond to security incidents.

SOC: Security Operation Centre, an in-house or outsourced team of cybersecurity professionals, that monitor the organizations IT environment for cybersecurity events.

Microservice: Microservice is a software architectural style where an application is developed as a collection of smaller, independent, and loosely coupled services, where each service is focused on a specific task and is able to communicate it to others.

Malware: Malware, short for malicious software, is any software designed with malicious intent to harm or exploit computer systems, networks, or users. Common types include viruses, worms, trojans, and spyware.

Phishing: Phishing is a cyber-attack where the attackers use deceptive emails, messages, or websites to trick individuals into revealing their credentials.

Microsoft Azure Active Directory: Microsoft Azure Active Directory is a cloud-based identity and access management service that provides authentication, authorization, and identity management. The service is currently undergoing a name change to Entra ID but will be referred to as Microsoft Azure Active Directory or AAD for short.

On-premise: On-premise is a traditional computing model where the organization owns and maintains the computing infrastructure and software.

Proxy: Proxy is an intermediary server that acts as a gateway between client devices and the end system.

RDP (Remote Desktop Protocol): RDP, Remote Desktop Protocol, is a proprietary protocol developed by Microsoft to allow users to connect and control remote computers or virtual machines over a network connection.

SSH (Secure Shell): SSH stands for Secure Shell protocol, which is used for secure communication remotely over an unsecured network.

Software as a Service: Software as a Service is a cloud computing service model where software applications are delivered over the internet. Users access a third-party provider's service, eliminating the need for the organization to manage and maintain the service.

1 Introduction

Today's business environment is complex and it will inevitably become more complex as new emerging technologies are added to the compounded collection of information technology systems. The drive towards more complexity will only grow as European Union has set a strategy to become the leader in data-driven technology. This data has to be protected in businesses. As business is no longer possible without various IT systems within an organization, the need to protect them only grows. The basic IT functions of businesses is to deliver and maintain productive and secure systems to create value through enabling other departments. Identities, authentications and authorizations must be managed in order to operate these IT systems in a secure manner, which is the main purpose of identity and access management. Identity and access management activities are considered mainly as preventative approach to security and controls, compared to detective or responsive methods, which both are usually part of operational information security. Preventative methods mainly enforce policies and counteract security breaches before their presence in the system. Detective and responsive methods fight to identify and alleviate a security breach that has or is already occurring. The contemporary preventative measures against a possible breach through the lens of IAM include: Strong and secured user authentication, up to date authorization, sufficient logging of access, timely de-provisioning and user credential management, and lastly governance and compliance on a continuous manner. [1]

This thesis was written in collaboration with target company, to create documentation for high availability and disaster recovery for a privileged access management tool that is being implemented and not yet fully functional. The work is part of the organizational identity and access management posture and was a relevant need for the target company.

The research sources for the thesis were searched for using online sources from Google Scholar and UTU Volter databases. In addition to these, the thesis will leverage available frameworks, within the intellectual property rights of the target company's own sources and existing protocols/frameworks in identity and access management. Thesis was written using an agile approach, having weekly checkpoint iterations with the assigned supervisors.

The main research questions to be answered by this thesis are:

- What are privileged access rights?
- Why should privileged access rights be secured?
- What type of protection mechanisms can and should be applied to them?
- How the resilience of these protection mechanisms is ensured in case of a disaster?

These research questions are to be answered in this thesis by using design science research method and following the guidelines of design science. The thesis will begin with a literature review of identity and access management and the resources accessed are to be no older than ten years old to ensure the information reflects modern IT infrastructure. However some of the principles and cited sources are older, but still reflect current environment due to their strength in adaptivity to modern information systems.

The thesis will continue by defining the problem scenario which design science artifact must be created for. After outlining the scenario, the artifact will be implemented and results will be evaluated according to design science principles.

The thesis will end with a conclusion of the research, where the research questions are explicitly answered.

1.1 Research Methodology

Information technology is the subject area for the thesis, which provides an opportunity to use design science research methodology. The methodology is used, when creating new organisational and technical procedures. Design science is a prescriptive research approach to connect design and actions. It mainly consists of researching a topic, prototyping a solution for a problem, and creating a design to partially or fully resolve the problem. Therefore, design science enables and provides the researcher with an opportunity for observing a phenomenon and creating improved designs resolving identified problems. Design science is inherently a problem solving process, which requires the creation of an innovative, and purposeful solution. The design science process is defined by the connections between research and the knowledge base of the research. The main focus is creating a connection from information in the existing knowledge base with the design science research methodology to purposeful artifacts, which serves a specific purpose. Usually and in this research, this is carried out in an organizational environment. [2] This thesis is done in

collaboration with a target company, which creates a possibility to use design science research for existing practical engineering problems at the target company.

1.2 Methodology Application

This thesis will be following the guidelines set by Hevner et al. [3] In their article on design science in information systems research, as it will provide a good framework for this thesis' research questions. There are seven guidelines outlined by the article, that describe the way of working in design science. First is to produce a viable artifact in the form of a construct, a model, a method, or an instantiation. The subject of this thesis is to produce a model for disaster recovery in privileged account management tool in an enterprise environment (for target company), which would satiate the first guideline. Secondly the problem relevance, the problem should be relevant to the business for the need of design science. The problem is very relevant for this thesis, as privileged account management is considered the most important aspect of IAM in the target company as well as in many other organizations. The third guideline is design evaluation, this thesis will state the utility, quality and efficacy of the designed artifact in a chapter that evaluates the completed work. Fourthly the design science research must provide clear and verifiable contributions in the areas of the design artifact. The fourth point can be considered achieved, as security guidelines in the target company requires that a service recovery plan is created and the research artifact contributes clearly to the required document. Fifth guideline is the research rigor, where it is stated that rigorous evaluation and construction methods must be used on the artifact. The guideline is met with the mechanisms used during the research process, where the artifact is created with an iterative process gaining insights through continuous feedback from the project team. The second last guideline states that the design science research is a search process utilizing available means for creating an purposeful and effective artifact. The research meets the guideline as the artifact is created into a new area, which cannot be done without using relevant and multiple means to create an effective artifact resolving or improving resilience and disaster recovery. Lastly the communication of the research. This thesis will be the academic background for the solution, and it will be presented at the end of the research for the target company's management. This will create a possibility for others to build advantages for further extension and evaluation of the completed artifact. The fundamental questions of the design research, such as "What utility does the new artifact provide?" and "What demonstrates that utility?" will be answered in the conclusion of this thesis.

1.3 Observational Work

Observations help the researcher to gain better understanding of the problem and the need for an artifact. The observations will be done by participating in team meetings and working within the target company's identity and access management team, reviewing the ways of working, and applying the best practices into the research from the IAM and project teams. The research will be evaluated weekly through an agile progress evaluation process. The observations will be done while working in the target company's privileged access management improvement project team. The observations are performed approximately during a half year period.

Observing the existing documentation and understanding differences between the target company and other organizations are key steps in the thesis. These steps are needed in order to create an effective solution for the target company that fits their needs. Solutions are dependent on the industry and the ways of working in an organization, since different industries have different needs and requirements.

2 Overview and Key Concepts in Identity and Access Management (IAM)

Foundations for identity and access management are based on a set of key concepts, and their context provides basis for the importance of identity and access management. This chapter begins with an overview of identity and access management and continues by describing the key concepts in identity and access management.

Frameworks, policies, technologies and processes create basis for the identity and access management. The identity and access management is used to manage and control access to assets and resources within an organisation's IT environment. [4] Main objective of IAM is to enable the right people to access the right resources at the right time. This minimizes risks of unauthorized accesses to sensitive data. With more organizations using data as a resource and industries evolving to be data driven, the importance of IAM as a security measure is highlighted. IAM has an integral role in modern enterprises' security foundations by protecting data and preventing data breaches as it allows organizations to limit their cyberattack surface. The identity and access management will have a key role in regulatory compliance. The GDRP requires already limiting access to personal identifiable information (PII) and the NIS2 directive increase requirements in identity and access management within the EU area. The thesis will revert back to topic of the regulatory compliance in the upcoming paragraphs and elaborate the topic further. Identity and access management can work separately as identity management and access controls. The identity and access management is usually bundled together as both of the elements are needed for accessing IT assets and resources. Access management makes sure that given resources are only accessible for users that have been granted the sufficient amount of access. Access management alone does not define identities which are used to access resources, but manage access rights associated with accounts linked to identities.

Identity and access management at application layer is often driven by requirements, for example from contracts and regulation, however identity and access management impacts also information security. IAM can directly reduce the risk to the biggest cyberthreat as of 2022, which was ransomware. Ransomware is a type of malware from crypto virology that encrypts the victims data to make it inaccessible in hopes that the victim would pay a set amount of money to release the data. Ransomware were 50% of Microsoft's cybersecurity recovery engagements in 2022 and ransomware incident engagements participated by

Microsoft revealed that insufficient controls over privileged accesses and lateral movement were root causes for a successful attack in 93% of the cases. [5] Similarly half of over 500 organizations in a survey by Fortinet had fallen as victim of a ransomware attack, which makes it the biggest threat within modern cyberspace. [6] By these metrics, it is can be assessed that privileged access management, which is one of the key parts of identity and access management is one of the most important controls to create robust cybersecurity within organizations. This notion was part of the risk assessment for the privileged access management solution in the target organization.

Identity and access management consists of three main categories, which are authentication, authorization and identity management, with all of them tying into identities. These three categories are divided by their own unique abilities and features for creating a robust IAM structure. As seen in Figure 1, IAM works as a middle control point to access target resources in the organization. The main objective of IAM is increasing an organizations cybersecurity posture by making them more resilient in case of an breach or attack. IAM can have a mitigating effect on affected resources and resist the attackers lateral movement by restricting access to resources within the organizations IT environment.

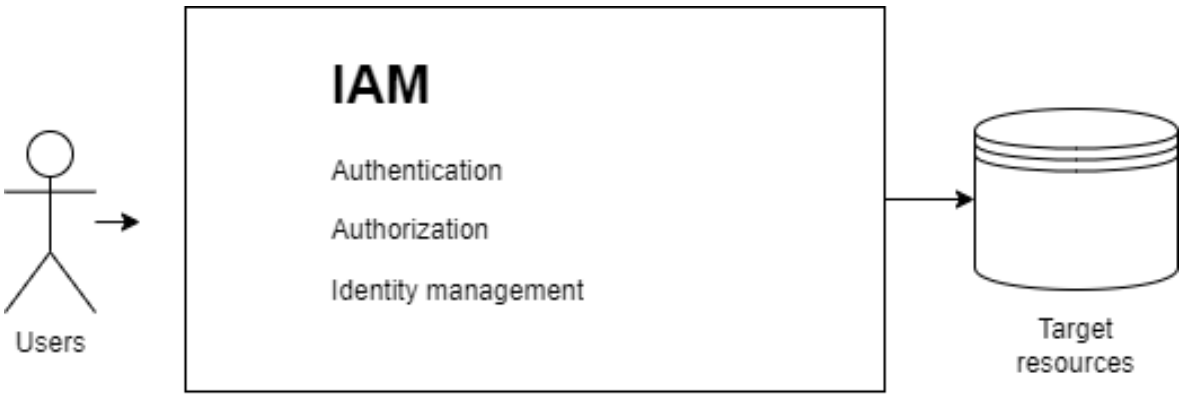

Figure 1, IAM position in IT infrastructure

Figure 1 shows a very simplified version of IAM's position as a proxy from users to targets, based on encyclopaedia of cloud computing's identity and access management functional architecture. [7]

2.1 Relevant IAM frameworks

There are multiple frameworks, standards and guidelines for an organization to use when designing and implementing identity and access management processes, controls and procedures. The frameworks, standards and guidelines supports also compliance in the

identity and access management. [8] The frameworks, standards and guidelines provide organizations with set of approaches, processes, procedures and attributes, which they need to adapt based on their threat landscape, requirements (compliance and business), risk appetite, risk management practices etc. to sufficiently control and manage the IT environment

A network information systems directive 2022/2555 (known as NIS2) is entering into force in the European Union. Replacing older cybersecurity directives, that command organizations working in EU member countries to increase their cybersecurity capabilities and national cybersecurity strategies. Therefore, NIS2 will increase the regulatory compliance requirements for cybersecurity governance, capabilities and management for organizations in EU member states, such as in Finland. The directive aims to build and ensure sufficient cybersecurity capabilities across the Union within organizations that provides electronic communications, digital services, food, critical product manufacturing, waste management, postal services, or public administration. A major part of cybersecurity capabilities are IAM processes and controls that will be affected by the legislation. The target company of this thesis falls under the original essential entities of the scope of the directive and therefore has to comply to the directive while operating within the borders of the European union. The non-compliance penalties for essential entities can be up to €10,000,000 or 2% of the total annual turnover in the previous fiscal year of the company, whichever was higher during the previous year. [8] This creates motivation and an incentive for organizations to be NIS2 compliant, to avoid penalties. The original NIS directive set into policy in 2016 was adopted in Finland 2018. The future directive that Finland will apply is expected to be more strict, which should create an uplifting and technologically advanced environment to Finland compared to the previous. [9]

One of the ways to prepare towards compliance is following the ISO27000 standard family frameworks produced by the International Organization for Standardization (ISO). [10] The standard published in 2005 with International Electrotechnical Commission (IEC) was revised in 2022 to create detailed requirements to establish, implement, maintain, and continually improve information security management systems. This creates security governance and management structures and practices for organizations to secure their information assets. The ISO 27000 standard family provides an approach to organize information security management system based on risk driven process, where threats, business impacts and risks are systematically identified and assessed. This enables a organization to design and implement security policies, processes and controls that mitigates

the significant threats and risks. These controls can be audited by a certification auditor, which tests the depth and extent of the implemented controls in an organization and can certify companies to be ISO certified. The ISO27002 standard from the ISO framework has more detailed requirements for information security including the IAM topics. It considers access controls, human resource security and physical and environmental security. Usually physical and human resource security are considered as part of a traditional security and it is up to the organization in question if they are more relevant to identity and access management or other functions. Number of the controls are associated with the identity and access management despite where the accountability and responsibility for the controls are in the organization.

National Institute of Standards and Technology (NIST) cybersecurity guidelines are another relevant set of the best practices. Specifically NIST 800- which provides a catalogue of security and privacy controls. The framework is mandatory for U.S federal government agencies, but is also relevant with sharing important best practises for EU members. NIST framework emphasises the importance of implementing a zero-trust architecture with a focus on resource protection of an organization. [11] Zero-trust principles include “never trust, always verify” and “least privilege” both being essential in identity and access management. [12] These ensure that authentication and authorization are performed before every stage of a digital transaction. The core thought behind these frameworks being that there is no automatic trust in any user inside or outside the network and access is therefore granted in a need-to-know basis. These frameworks lay the groundwork for improvement in identity and access management security posture.

2.2 Identities

Understanding identity and access management begins from understanding what identity means, what types of identities there are, and what are the important aspects of identities. Identity can be a multi-faceted abstract concept, but this thesis will strictly focus on digital identities and their technological aspects. A digital identity is the sum of attributes and information that establishes and differentiates an individual user or an entity digitally. It is comprised of personal, demographic and possibly user input data that uniquely identifies an individual. This forms a basis for a link between the physical and digital worlds. [13] “An entity may have several identities and human entities almost always have many identities.” (Gupta. Et al, 2012) Therefore, a system is needed to manage these identities, such system is

called an identity management system (IDM). IDM's handle the provisioning, upkeep, and deprovisioning of identities. [14] These tools are sometimes commercially called Identity Governance & Administration tools (IGA's) and they are usually branded as such, if they provide additional features to traditional IDM's.

Digital identities can be defined with attributes of an entity, usually these come from the human behind the identity. Digital identities are not limited to humans entities only, but can be given for example to machines, servers, and service accounts. Within the scope of the thesis, only digital identities of humans are taken into consideration. A persons digital identity can be defined using attributes. These attributes are usually natural attributes of a person such as name or gender, but can also be numerical attributes such as phone number or address. The natural attributes of a person can be verified by biometric data since they can uniquely define individual people. Biometric data traditionally is fingerprints, iris texture, or facial structure, but new emerging biometric technologies could be used in the future. Digital identities are created by capturing identifiable attributes, storing the attributes and tying them to credentials, which can be used to authenticate and access electronic services. The method of capturing the attributes directly impacts the confidence of the information. NIST defines Identity Assurance Levels, that convey the degree of confidence that a persons claimed identity is their real identity, with level one being lower confidence and level three being high confidence. [15] The identity assurance levels are handled by the security in the target organization, and all work with and within the privileged access management tool is considered to be high impact or assurance level three work.

Digital identities are stored in an electronic database, where they can be used to validate and authenticate users in the organizations IT environment. The challenges of a digital identity is that it is used to store personally identifiable data, which has regulatory implications within the European union. According to the GDPR, names or location data is considered personally identifiable information which can be used to identify an individual. [16] This means that the stored information has to be compliant to data protection rules that enable the privacy of the users. If an organization processes data for the sole purpose of identifying someone, then by the definition it is personal data.

2.3 Authentication

Authentication is the process of verifying the users attempting to access a system or a resource. This step ensures that only authorized individuals or entities can gain access to

protected information or services. Authentication mechanisms provide the support for proving the individual is who they claim to be, thus establishing trust and enabling authorization. There are multiple ways of authenticating an user or identity and the following ones are an example of ways to authenticate entities within the scope of this thesis. Authentication is based on four authentication types: What you know, what you have, what you are, or where you are. These are simplifications of cognitive, possessive, physical attributes or locations. All of these methods can be used to authenticate an user. [17]

2.3.1 Password Based Authentication

It is possible to assume that all modern IT services use some way of authenticating users using passwords. [13] Passwords are usually considered to be a key that unlocks a lock to reveal something secret, traditionally protecting access to something of greater consequence. [18] At their most basic, password are a way to tie secrets based on identity, even outside the context of information technology. The basics of modern password authentication is that the user provides an unique password that corresponds to their account or identity. This password is then compared to a stored version of it, typically in hashed form using cryptographical algorithms. Then access is granted to the resources if the password matches the hash. [13] Password authentication can be made more secure applying password complexity requirements, usually over ten characters to make the time required to brute force the password long enough to not be worth it for the attacker. Assuming there are no other glaring weaknesses in the storage of the password, such as storing them in plain text, or in an unsecure location. There are other ways to implement mechanisms to detect and prevent brute-force attacks, but none are as effective as creating longer passwords. Fundamentally the problem with password authentication is that they can be shared with others or be impersonated and no technical cybersecurity control can manage this risk.

Passwords have a long history of authentication use in information technology and passwords and PINs are well known and trusted authentication systems to users. This also provides the fundamental problem with passwords, that their weaknesses are also well known. Attackers can use guessing attacks and brute-force to falsely authenticate themselves in order to do an identity theft attack. Passwords should also not be shared between systems, always creating a new one for a new system. New passwords create problems for users as users usually create more vulnerable passwords or share passwords between systems because of the reduced usability of unique passwords. The obscurity and uniqueness of passwords is

generally held as the most important aspect of creating a strong password, however with the rise of stronger and faster computing power the length of the password has become an ever-growing aspect of password safety as shorter passwords are more vulnerable to the already mentioned brute-force attacks. [19]

2.3.2 Multi-Factor Authentication (MFA)

Multi-factor authentication, sometimes reduced as two-factor authentication is an authentication method used to establish an identity using multiple different authentication methods. Multifactor authentication is only applicable if multiple authentication methods such as passwords and biometric authentication are used in conjunction to create more trust in the individuals identity. [13]

Multifactor authentication significantly mitigates the risks associated with only single-factor authentication since if one authentication method is compromised, the attacker would need to overcome the additional layer of authentication to access the resources, greatly improving the security of the system. It is important to recognize that separate categories of credentials are used to improve the security of the authentication, meaning more than one method of authenticating the user is needed. Multifactor authentication can be used as a good tool to establish information protection with advanced or more critical levels of information being restricted with multifactor authentication. [20]

2.3.3 Biometric Authentication

Biometric authentication is the use of human biological attributes to verify the identity of the user. Biometric features of humans are rarely changing therefore, they provide a good basis of trust for authenticating an user. Biometrics can be categorized differently as physiological and behavioural biometrics. Physiological biometrics are related to human bodily features, while behavioural biometrics are related to certain kind of behavioural of an individual. Behavioural biometrics require much more investment to apply since they require long tracking of user behaviour which usually raises privacy concerns in users. [13]

A biometric authentication generally extracts a feature of an user that can be matched, for example a fingerprint. The fingerprint is then compared to stored fingerprints of the given user, usually taken earlier to create a database of authorized users, and if a match is found the

user is authenticated to access the resources limited by this system. Biometrics are widely used with government ID cards and passports.

Physical biometric authentication can be subcategorized to head and hand authentication. Head biometrics include features found on human heads, such as face in facial recognition, ears with audio authentication, or eyes in iris scanning. Hand biometrics include fingerprints, palm prints and hand geometry. [13]

Biometrics raise other privacy concerns as well since they rely on highly personally identifiable information, and therefore risk of personally identifiable information leakage is higher. The concern with storing personally identifiable information ties back to GDPR and the need to comply with regulations. A fear of users' biometric data being used to potentially discover other element unrelated to identity verification might also be a factor to hinder the application of biometric authentication. [21]

Limitations of biometric authentication are the inaccuracy of the captured data. The scanners used to capture biometric data are not fully accurate and can add noise to the data which might result into incorrectly matching individuals in the database. Distinctiveness of humans is not always crystal clear, for example identical twins have been a problem for facial recognition, but they are not impossible to identify from each other. Adding more parameters to the facial recognition can improve the results, but adding features always comes with additional costs. Another limitation of some biometric authentication is missing features of the humans they are applied to, some people might be missing fingers or eyes needed for the authentication therefore, another method of biometric authentication is needed. Plastic surgery might result in the authentication to be mismatched or possibly even spoofed if done accurately enough. [13]

Biometric authentication offers several advantages, making it a compelling choice for enhancing security in the future application. Biometric authentication provides a higher level of security assurance, since biometric data is more difficult to replicate or forge. Added benefit of choosing biometric authentication enhance user verification with multifactor authentication is that it is considered to be easier for users, eliminating complex passwords and replacing them with quick scans of user biometrics. [21]

2.3.4 Single Sign-On (SSO)

Single sign-on allows the user to authenticate into multiple different applications using the same credentials within an organization. The main benefit of using single sign-on is to reduce the amount of different authentication when switching applications. Single sign-on is heavily tied to having appropriate identities in place since it eliminates having different credentials for the same identity. Having a user only remembering the credentials to their identity reduces the compounding risk of having multiple credentials which increases the attack surface for illegitimate authentications. Other than that the reduced amount of service requests for forgotten passwords is a quantifiable benefit that can be seen immediately after adopting single sign-on features. [22]

Advantages of SSO are more on the user experience end, since it creates some problems within governance of the identities. The identities can have conflicting data when they are tied to simple attributes, such as a name or job title. Therefore, a robust and a detailed global governance model in the identity management system is needed for the user identities. [13] Other major consideration from a cybersecurity standpoint is that the SSO could create a single point of failure. If the SSO credentials are to be compromised it allows access to every other application that uses the single sign-on feature. Hence, it must be secured with the strictest cybersecurity controls, or apply separation of duties by having critical targets out of the single sign-on scope. If enough trust is built over this single point of failure it should still be more secure than the user having multiple different credentials as this grows the attack surface and increases the possibility of hazardous password management. Many state of the art IAM solutions provide safe credential management, and thus the application of single sign-on feature should be weighted individually within an organization.

Single sign on cannot be considered a fail-safe mechanism to avoid identity theft as there are multiple way to bypass authentication using single sign-on vulnerabilities. [23] It is possible for an attacker to impersonate, or steal cookies by eavesdropping the in-browser communication that store single sign-on information to bypass authentication to target sites.

2.4 Authorization

Authorization is one of the important aspects of IAM together with authentication. Where authentication verifies the user, so the user is the identity they are tied to, authorization is granting the necessary actions and resources the identity is allowed to access. By defining and

enforcing access control policies, authorization ensures that only the authorized identity can perform the required activities. It is important to note that authorization must be strict to only to the required resources to limit the access in an IT environment. The purpose of authorization is not only to grant access privileges to users, but also to deny access of those that do not fill the requirements of said resources. [13] This ensures the principle of least privilege.

Information protection is the systematic approach for identifying, classifying and protecting information consistently across an organization. Resources are generally limited by data classification in order to assess the criticality of the information to ensure critical resources are protected. Authorization plays the role of managing access to these resources according to the needs of the users. It is considered best practise to ensure the principle of least privilege for users in an organizational information system. Principle of least privilege is as the statement implies, a control to give as little privilege to user as they need. This limits the users access to resources they should not need in their daily work, making the resources more secure by limiting access in case of a breach on the users identity. Limiting access can be done with the help of ISO 27002 framework. [24] ISO 27002:5.12 provides a basis for data classification that has been used widely in different organizations, including the US armed forces. The control should not be taken as is, but fitted to suit the organizational needs, giving a good basis to create, to each their own levels of confidentiality, to later enforce authorization over.

Generally information confidentiality is classified into four categories (unrestricted/public/available for all, confidential, secret and top secret) for example by NIST and ISO. The information classification categories can be used to define requirements for management of information security and identity and access management. The confidentiality classification of the information is determined by its sensitivity and the potential business impact of unauthorized disclosure of the material. Varying levels can be added in case of a need in an organization as the classification categories should not be one size fits all solution.

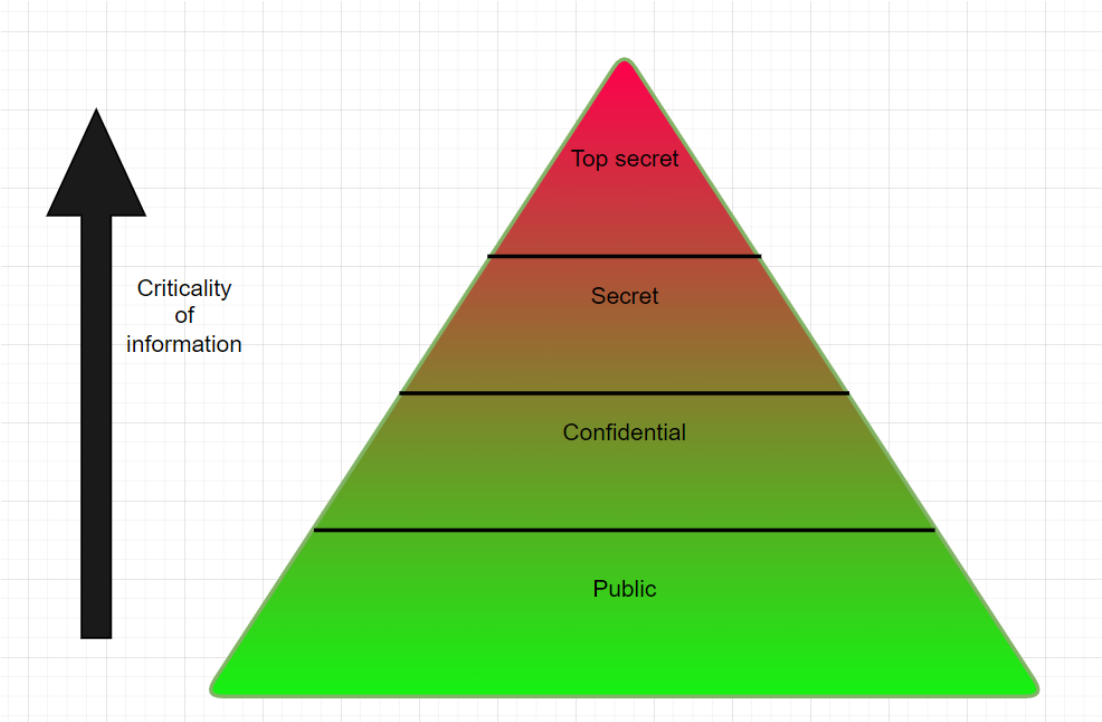

Figure 2, A pyramid illustration of common levels of information confidentiality classification

Confidentiality classification levels should follow a stricter set of access control the higher up the classification is. Less amount of users are authorized the stricter the classification gets with zero trust principle in mind, the user should not be able to see information above the levels required. By the confidentiality classification categories in Figure 2, it is possible to enforce authorization over information in an organization:

Public information being something that has no impact on the business and can be shared among outsiders of the organization. The model might not be relevant to apply for viewing the information, but creating it.

Confidential being information of low sensitivity and general business nature, for example intranet news in an organization.

Secret being sensitive and restrictive in nature, available only for a more limited group of business need-to-know, this is information that might have negative impact on the business being available for outsiders.

Top secret, where information is highly sensitive and may cause catastrophic negative impact on the organization, for example detailed plans of the organizational security. The information

that is classified with these levels can and should be authorized to be accessed using one of the different authorization models.

This way is not a definitive authorization level specifically in the target organizations, since it needs more details based on the information's location. For example the legislation on the use of personal data is very different in the EU because of the GDPR and might differ in EU countries, where some type of hybrid between confidential and secret is needed, based on the underlying information.

In the target organization, the confidentiality of the information are based on the legislation and contractual agreements between parties. A valid purpose for access, based on the information classification and the owner of the information is required to grant any level of authorization on the information.

2.4.1 Role-Based Access Control (RBAC)

Role-Based Access Control is a widely adopted authorization model that simplifies access management by associating permissions with clear predefined roles, offering a structured approach to access control. Managing individual access rights creates a big jungle of permissions with no clear visibility to the risks associated with privileged permissions, RBAC creates common roles that are more easily managed. Organizations can manage access rights based on job functions and responsibilities if the roles are well defined to personnel. For example the roles could include "Manager", "Employee", or "Administrator". With Role-based access control the roles can be easily understood by their names, which define the permissions, enabling a non-expert person to assign roles to personnel. Role naming simplifies the management of the roles for the administrators of the identities, but also creates visibility to attackers to target specific types elevated privilege accounts. To apply least privilege principle, the roles should have individual rights that are not shared between roles, to avoid creating a "Superuser" that inherits every aspect of these defined roles. Once the roles are defined, permissions are assigned to each role, that specify what actions and resources the specific role is authorized to do. By mapping these user roles access management becomes more scalable and manageable. The data classification hierarchy levels can be applied to access rights with RBAC. These accesses are mapped with the levels of confidentiality mentioned in the previous chapter, creating a hierarchy with increasing access controls to restrict users in complex organizational structures. [25] RBAC comes with the

downside of being slightly outdated as of 2023 since it cannot handle Realtime attributes, such as time of day and location.

2.4.2 Attribute-Based Access Control (ABAC)

Attribute-based access control is an authorization model that takes into account multiple attributes of the user identities and environmental conditions to make access control decisions. ABAC relies on policies that define rules based on user attributes that can be, for example role, department, clearance level. Attribute-based access control is sometimes referred to as Policy-Based Access Control (PBAC) for this element. In addition to these it also defines based on the resource attributes e.g., sensitivity or classification, which form the basis of the need for authorization. ABAC therefore offers a dynamic and context-aware access control, allowing more complex authorization policies based on more specific conditions and requirements than regular RBAC. With ABAC the need to define and engineer roles is unnecessary, only if the roles are not used as attributes, rather dynamically changing attributes should be used like location and time of day to determine the access control.

The problem arises with attribute-based access control when the permissions need to be managed. If RBAC allows ease of management by non-professionals by clearly defined roles, ABAC does the opposite of needing highly engineered attributes that might reach a number of thousands on some users. This creates a giant network of attributes making it very hard to manage, since it doesn't leave a clear audit trail.

Therefore, a combination of RBAC and ABAC is usually the feasible option to make a comprehensive information model for access control. Adapting a hybrid of both models enables balancing the positive effects of both models. [26] Balancing the positive effects is one of the desired effects of a comprehensive authorization model at the target company and therefore is applied.

2.5 Identity Lifecycle Management

Lifecycle management is the process where creating, deleting, management, entitlement change, and policy compliance are performed. [27] Individuals identity is managed starting from the creation to where the relationship with the organization is concluded. [28] Everyone and every machine that has access to the organizations resources should have an identity that is managed and has a person responsible for it, to avoid orphaned accounts which can be used

to attack the organization. The individual identities are managed by an Identity Management Software (IdM), which is capable of performing functions like, administration, discovery, maintenance, policy enforcement, and authentication. With these capabilities centred in the IdM the end applications do not have the need for their own identity storages and management capabilities, and the workload for the end applications is greatly reduced. Therefore, IdMs greatly simplify the management of large-scale distributed systems and help with the IT scalability of the organization allowing simplified applications for specialised tasks. [29]

The operational areas of identity management software include authentication management and authorization management. These operational areas serve as the base of operations for all other applications in the organization creating a firewall type structure, but for access controls in an organizations IT infrastructure. All of the access requests go through the IdM and are either allowed or disallowed to continue to the end destination after a policy check on the user identity. If the authenticated user fulfils the requirements, they are authorized to reach the desired resource by the IdM.

3 Privileged access management (PAM)

3.1 Privileged Access

To make sense of privileged accounts it is helpful to know the most common form of privileged access. The most common type of privileged access is within a user directory of an organization. Computers, users, and other components of an enterprise network infrastructure need to be managed to ensure their privacy and integrity. For this purpose directory services are essential to ensure authorization of users to end resources. These directory services must be managed by elevated privileges to ensure that regular users do not make changes to resources that might have an impact on more than intended. Regular users are usually only authorized to view certain content in the system, and more rarely make permanent changes into it, which ensures the integrity of the directory. Changes can result in accidental errors that might be harmful for productivity of and organization or even impact availability of their IT systems. Only risk is not accidental changes, but also changes from intentional attacks leveraging regular user accounts. Therefore, regular user accounts are not usually permitted to make changes into critical systems, which increase security of organization's IT environment. The role of an administrator is a job role with a need for privileged accesses, which enables them to perform their job responsibilities and make changes for example to the used directory services. Administrators have privileged access to create, delete and maintain the directory, and with the access only their designated account is eligible to make these changes in the directory. [30] If the correct implementations from the RBAC-model are in place, the regular user account breaches should not reach the same level of impact that a breach on a privileged access user will have.

Enterprises with a wide IT infrastructure, the most appropriate action to be taken is to delegate duties of administration to multiple IT personnel. The result of the delegations is that many administrators will have full authorization to multiple resources that they might not be eligible to see from a security standpoint. A fundamental problem for cybersecurity is users having too much visibility in an IT environment: The confidentiality of information is compromised if a persons without valid reason is able to see it. This also creates unnecessary attack surface. Attack surface is anything that can be leveraged in a cyberattack, in this case it would be the amount of elevated privileges. [31] The elevated privileges create an optimal point of entry for an cyberattack. The elevated privileged may give full access to any information in the breached system if an attack is successful. Therefore, it justifies the need

for a higher security measures on privileged accesses. The privileged access might not be always associated with an account of a human administrator. Privileged access can be granted for machines or service accounts that might work automatically. It is essential to keep these types of privileged accesses secured and activities logged for monitoring, and post-incident analysis, as there is no human monitoring activities performed with these accounts.

3.2 Threat Landscape of Privileged Access

Due to the elevated ability to do harmful behaviour with privileged access, it is important to understand why it is critical to protect these preventatively. A review of organizational and technical environments is needed to identify and define needs for a privileged access management solution. An analysis of the threat landscape for privileged accesses is also required. The following threat landscape is made based on the target company being an enterprise in the energy sector in Finland. [32] Energy sector is identified by European Union as critical infrastructure as disruptions in the energy sector will likely have effects to other critical sectors. Therefore, the processes in the organizations must be up to the standards set by the legislation. This includes the privileged access management processes and capabilities. Operations of modern organizations are relying on information systems and the energy sector is not an exception. Information systems are linked either through processes or technical dependencies. Therefore, understanding of dependencies is required when assessing criticality of individual information systems for the critical production operations. Additionally, risks associated with privileged accesses is not only driven by criticality of information systems. The threat landscape varies depending on characteristics of an information system and the related dependencies. Therefore, it is good to understand what threats and risks being in the energy sector apply. [33]

3.2.1 Physical and Political Threats

The main physical and political threats might not seem relevant as they are broad in scope, but might still have an effect on privileged access management. As the availability of energy is increasingly important for industries and consumers the availability of the systems facilitating the production of energy, in this case the availability of the privileged access management procedure, might have an compounding negative impact in a larger scale than one could expect. [34]

Starting from the largest picture available, global geopolitics as of 2023 have impacted energy prices globally, which impacts the availability of all electronic systems.

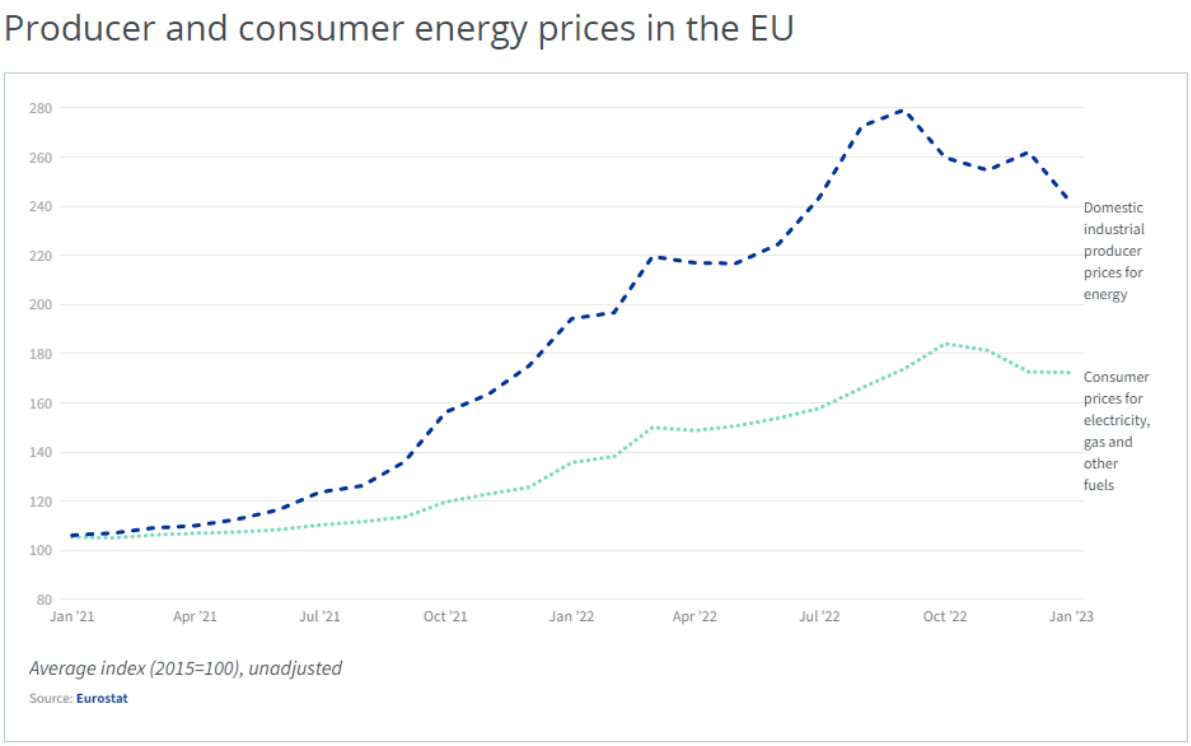

Figure 3, Electricity price hike in the EU [35]

Figure 3, shows the unprecedented rise in electricity prices in the recent years and for example the recent volatility in energy prices might affect the target company's policies for investing in cybersecurity, since it's hard to predict the gains or losses that might impact the company in the near future. Another uncertainty is the electricity availability of the vendor datacentres and servers hosting IT environment and systems, which has to be taken into account when considering the resilience of the system to outages or disruptions in the Nordic national grids or cross-country connections.

From a political perspective, the big threats to any electronic systems is the rise of hybrid warfare. [36] Hybrid warfare is a modern term for non-military instruments causing damage to utilities in an opposing country. Hybrid warfare is not considered as a traditional declaration of war and is continually happening in the cyberspace. It is mostly driven by state funded advanced persistent threats, which are introduced better in the cyber threats section. [34] The Finnish government has outlined that having information systems hosted abroad or in cloud systems, the threat of political and physical pressure is increased. Human error and sabotage/terrorism could possibly create large outages for the whole energy infrastructure in

Finland. Therefore, the IT-systems managing it or the production of energy should be hosted and managed in Finland if possible, although this is only advised and not mandated. Having independence over the management of IT systems through Finland creates resilience towards global political pressure and crisis situations.

3.2.2 Cyberthreats

Cyberattacks are a constant threat to any system in the digital space. Their objectives are to steal sensitive information leading to a loss of privacy or disrupt the digital system to cause adverse effects. The cyberattack perpetrators can range from individuals looking for profit, to state-funded adversaries looking to disrupt digital systems of another nation to gain influence through hybrid warfare. Detection and most importantly prevention is the key task for cybers defence of an organization. Continuing the top down view from global threats to smaller individual threats, the biggest concern of an critical infrastructure organization is an Advanced Persistent Threat (APT). An APT is usually politically or economically motivated large scale cyberattack with operators that have a vast spectrum of possible advanced targeting methods, tools, and techniques to reach and compromise their target organization. APTs are highly coordinated with a main objective to reach and maintain long term access in the target systems. Utilizing multiple possible intrusion methods APTs gain access to targeted systems and use lateral movement to access critical targets. Figure 4 shows the full lifecycle of an APT, but the focus with priviliged access is mainly in the third aspect, lateral movement.

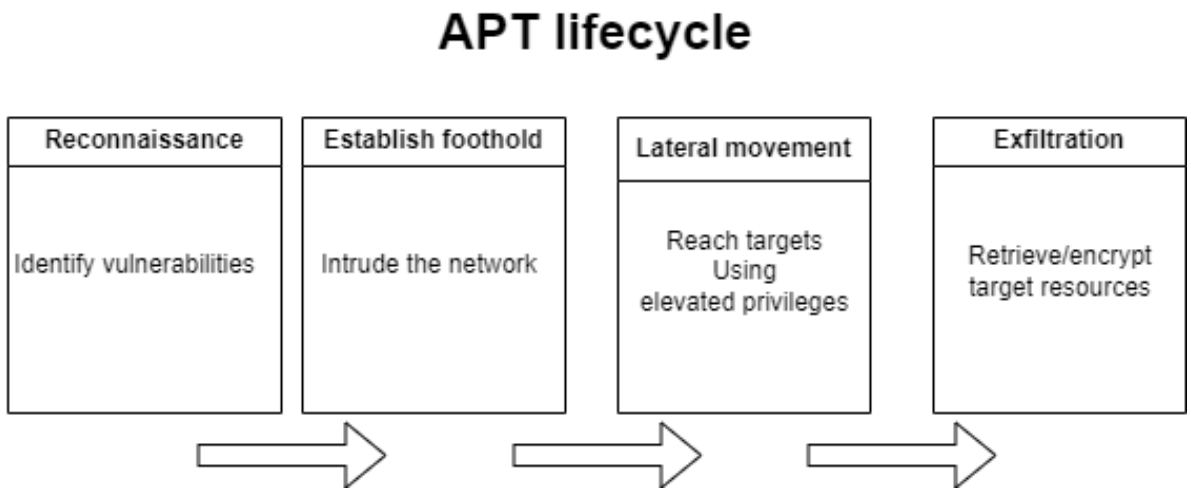

Figure 4, simplified workflow of an advanced persistent threat.

Crowdstrike, one of the leading organizations in network defence, describes lateral movement as the attempt to move from device to device or elevate privileges in a closed network, such as an organizations Domain infrastructure, gaining additional points of control in the compromised network. [37] These positions then enable the persistence part of APTs to stay connected to the target systems. Lateral movement relies on escalation of privileges, but if all privileged access are centrally managed, it sequentially reduces the possibility of the adversaries getting access to privileged accounts. If user identities are verified multiple times during escalation of privileges, for example logon to domain and PAM with multifactor authentication, the attackers have to gain access or bypass all the steps in the authentication process, which creates new layers of security to the privileged access. Additionally, lateral movement actions often leave behind malware that creates internal users onto the target system to possibly create seemingly authentic new user identity requests in the target networks information systems service management system (ITSM). Therefore the maintaining and provisioning of new accesses should be monitored and taken into consideration at an organizational level. An intrusion is very hard to detect if the attacker has gained elevated privileges during the attack. The network traffic can appear normal and the attackers often use built in tools of the environment to continue the attack.

Even regular user accounts can provide an entry point to gain access to any system, if managed incorrectly. Regular user accounts are user identities that do not have elevated privileges in critical systems, but are more prone to user error. Typically users may use weak passwords as stated in an earlier chapter or fall for phishing attempts. Phished accounts are usually gathered into big identity dumps that are then sold in the darknet for other attackers to use possibly during an APT. [38]

Ransomware is another relevant type of malware used in cyber-attacks. Ransomware were the biggest threat to organizations in 2022 according to the Microsoft Digital Defence Report. [5] Ransomware is a type of malicious software designed to encrypt the victim’s files or entire computer data, rendering it inaccessible until a ransom is paid to the attacker. The attackers demand payment, often in cryptocurrency, in exchange for a decryption key that should unlock the victim’s files, although it cannot be guaranteed as the attacker does not benefit from it. The target of ransomware is usually a large, high-value organization and the victims are chosen based on their ability to pay the ransom to the attackers. [39] Attackers mainly focused on large organizations are called Big Game Hunters by Crowdstrike, implying the connection in their habits to traditional wildlife hunting. These big game hunters are

highly organized and utilize modern ransomware in a similar way to a software-as-a-service (SaaS) business model by using ransomware-as-a-service (RaaS). The big game hunter groups are likely targeting the target company based on its wealth and size. Therefore, these groups are relevant for cyber security threat modelling of the target company. Impacts of a ransomware attack to an organization with complex IT environments are difficult to identify and evaluate. However, it is not only a direct financial loss, but can have a compounding added loss if the target systems expose sensitive customer data, which is very likely while operating with customers that have to give payment information and location information to receive services from the company. In addition to the financial losses, the ransomware will create an operational disruption to critical services in the organization possibly causing the loss of thousands of workhours for the workers in the company. Lastly the reputational damage might cause the whole organizations brand to devalue, and long term reputational damage is very hard to overcome. Clearly making the choice to not rely on paying a ransom, but preventing the ransomware attack from occurring is likely a better option to operate under the circumstances given. For example a ransomware attack on the shipping and logistics group Maersk, led to them reporting a 200-300 million dollar loss. [40] Ransomware is constantly pressing and evolving problem for organizations and therefore, designing and implementing sufficient preventative and management measures based on associated risks is needed.

There are various cyberthreats on privileged access, however the APT and ransomware are the most important ones to prevent. They can be considered umbrella threats to privileged access as APTs and ransomware attacks usually have to use multiple different strategies and vulnerabilities to have an effect on the system. It is important to recognize that not all problems addressed by identity and access management are big threats from outside actors, but help to keep all identity usage streamlined within an organization. This helps to prevent an inside actor to make accidental or intentional harmful changes to critical IT systems.

3.2.3 Ordinary Threats

Regular threats to privileged access management are usually things that might compromise the service. Third-party vulnerabilities, physical security threats, configuration errors, denial of service through physical means. All of these can be considered threats that are hard to address when designing a service that relies on other vendors for production. Although not

unique threats to privileged access, the setting of privileged access can create scenarios which have compounding negative effects on other services.

3.3 Privileged Access Management Solution

Gartner, one the leading international IT research and consulting organization, states that a privileged access management tool manages and protects accounts, credentials and commands that offer an elevated level of technical access to IT system. [41] These PAM tools can be software, software-as-a-service or hardware appliances. Gartner states that there are three distinct approaches to privileged access management. Firstly privileged account and session management, which is the managing of passwords and other credentials for privileged accounts, meaning the passwords are changed at definable intervals or in the occurrence of specific events. It is important to note that password management is not comparable to session management as theoretical once authenticated session can be active indefinitely, unless some type of session management is applied. Secondly privilege elevation and delegation management, which ensures that specific privileges are granted on the managed system. Lastly secret management, which is often used in agile environments where it manages and stores credentials for individuals or multiple users. Key features of a PAM solution are restricting accesses to shared secrets and rotating the passwords of accounts after one user has checked out. Mainly securing sessions and having session management are key points of interest from an IAM perspective.

The target company is seeking improvements to the following capabilities and benefits from a PAM tool. The PAM tool should provide just-in-time access to critical resources, where privileged access is not granted for longer than needed for the actions to ensure least privilege principle. In addition to just-in-time provisioning, it would be essential for the PAM tool to provide a secure way to remote access critical targets using encrypted gateways and record the given remote accesses for compliance auditing purposes. Lastly, the process must be automated, there is no need for a tool that creates more work. One of the goals of implementing a PAM solution is to reduce the workload that privileged access creates. Features seen beneficial, but not a necessity, are the discovery of accounts with privileged access in the target environments in our network, monitoring of privileged access events and an additional layer of multifactor authentication to the target servers. The choice of a privileged access management tool will have an effect on the end artifact, security and possible vulnerabilities limit naming the solution outright. However, the currently chosen

PAM tool with majority of the listed capabilities under evaluation for an implementation.. The disaster recovery scenarios will be identified based on the PAM environment architecture. Due to timing of the project, the scenarios are identified and evaluated based on the current architecture.

With the chosen tool enabled, the end result should ideally provide full visibility to all privileged accounts and identities, human or service. With the full visibility it would be possible then to apply least privilege principle to resources considered critical to production. Governing and controlling the privileged accesses is more robust and an audit trail is left to ensure implementation of the requirements, in example processes, controls and other relevant elements. The target resources of the chosen tool can vary. The usual targets of the PAM tool are on-premise, such as physical servers, virtual servers, network appliances, and data centres that host the target resources. On-premise targets are installed in physical premises of the company rather than a remote facility. On-premise targets differ on a network level from cloud targets, since they can be accessed directly and are usually more critical to the organization. Cloud targets include for example AWS, Azure or Google Cloud subscriptions, instances and other resources. They are the big cloud service providers that enable a remote access from the target organization to their services. Other cloud targets can be business applications like Salesforce, SAP, Workday or privileged applications like GitHub and GitLab for software development. Cloud services provide ease of use for the end users, but require an additional layer of security as they are not managed internally. Lastly, the target of the tool can be databases, which provide the data for all of the above targets. With all of the above targets, it is certain that NIS2 regulation will require strict privileged access management controls and it is crucial that the tool will provide a possibility to have controls that enable the regulatory compliance needed for the future legislation.

The tool chosen will specifically help with three key areas within the target company's IAM development schema. It will significantly improve employee experience by easing the way to handle session management with administrative work by reducing the interfaces an administrator has to use. It will enable single-sign-on to on-premise targets previously unable to have the functionality and provide multifactor authentication to targets, overall improving authentication. The tool will allow stricter following of the zero trust principle, by enabling password-less logging in to targets, reducing attack surface and limiting visibility to key resources covered by the PAM tool. Furthermore, the tools ability to manage sessions making

sure there are no standing elevated privileges. The tool will be an additional layer of defence between the users and target systems.

The key features of privileged session management are ability to approve privileged access for a single session based on associated IT service management tasks, grant privileged access rights only for duration of a session, limit length of sessions, terminate active sessions and provide additional limitations to sessions (blacklist commands etc.). In almost any IT environment, there are systems with non-personal accounts with highest access privileges. This weakens audit trail and complicates monitoring use of these accounts. With privileged session management the access is immediately cut off or outright not granted if the identity is not authenticated properly. Basis for robust privileged session management is hiding secrets used to authenticate sessions from users. Additionally, secret associated with a single session are immediately changed after authentication or temporary. This combined with proper authentication methods, such as biometric authentication and monitoring of behavioural data will create a secure session for the administrator to work. These session are monitored to find anomalous or risky behaviour to either alert needed personnel or to trigger an automated suspension of the session.

3.4 PAM Architecture

For the privileged access management tool to work, it has to rely on underlying server infrastructure. The term infrastructure refers to supporting hardware and software for a service to run, which can include servers, data-centres, to web servers and operating systems. [42] The tool consists of multiple microservices that work together producing the end product. The key feature of the tool is called an application instance, which relies on data-centres hosting servers and web servers, on some type of operating system, to reroute connections to target systems. The main additional infrastructure features of the tool are load balancers, SQL-based databases, carriers, and web proxies. Key enabler of the tool's features is depended on a centralized authentication service, which is relied on to authenticate users to their identities using either strong credentials or biometric authentication. Creating an independent multifactor authentication based on biometrics is hard and expensive. To implement this feature and maintain it, is considered to be out of scope for any project in the organization.

Load balancing of the applications traffic will be done through the load balancer. A hardware load balancer is a hardware device with a specialized operating system and a software load balancer is software within existing hardware that handles the traffic towards

the application instances. [43] The solution for load balancing cannot be described to be either of these traditional categories of load balancers, however the technical details of the load balancers operation is not crucial when designing the PAM solution.

The PAM tool will be hosted by multiple application instances, which redirect end user RDP and SSH traffic towards their desired targets in other protected systems. The end user will make contact directly to a load balancer service that distributes network traffic equally across these application instances. After which the application instances handle the traffic and send it towards the databases or target servers of the end users depending on the type of the request. Resiliency and service recovery must be consider in all of the architecture decision to ensure availability and recoverability, which are aligned with the requirements. The application servers will be ran as individual instances of the application, with them being capable of working independently of other instances of the application. As seen in Figure 5, the route to the targeted systems might vary depending on the connection. Web access through HTTPS needs two additional parts of infrastructure to create a secure connection to the protected resources.

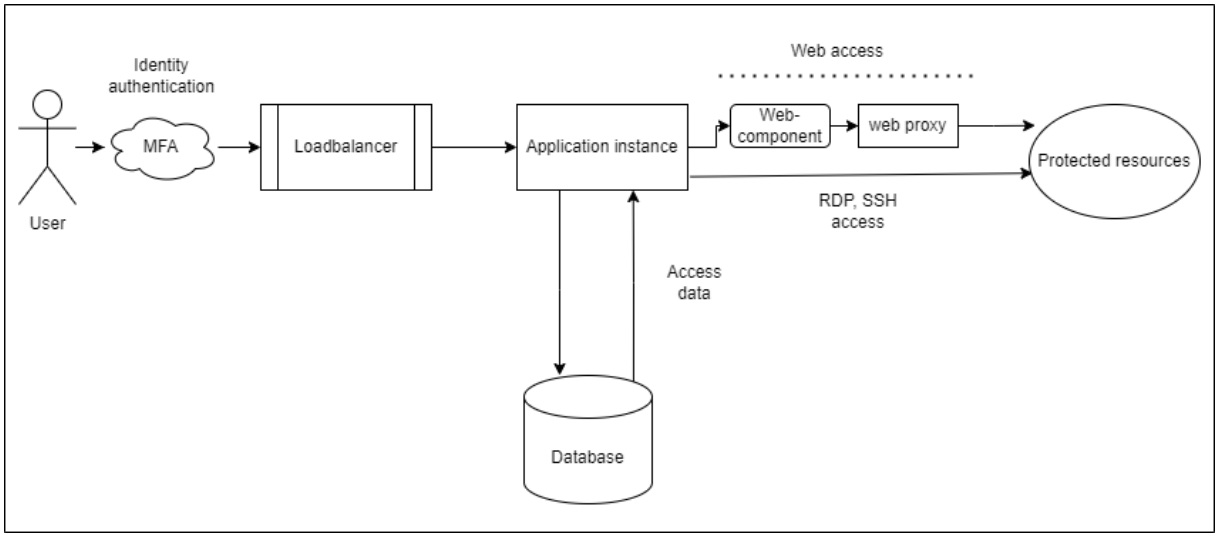

Figure 5, PAM tool architecture

The PAM tool will work as a gateway or a buffer structure with remote connections, usually secure shell protocol and remote desktop connections between users and target servers. Logical integration of the privileged access management tool will take user identity information from the inhouse user repository, which is part of the Configuration Management DataBase (CMDB). The logical integration of the PAM tool is visualized in figure 6. The user will be authenticated in the directory services where the data flows to the PAM application.

The privileged access management application can then create, modify, or delete the user's access towards the target servers. These changes are sent to the ITSM and CMDB tool that will update the usage of said identity and every change made in the identity will be stored in the privileged access management tool's own database. The privileged access management application queries access rights from the CMDB and is used to fulfil the requests made in the company's ITSM tool. This creates dependencies on a few outside sources of the PAM application for it to work properly, which have to be accounted for when considering its business continuity.

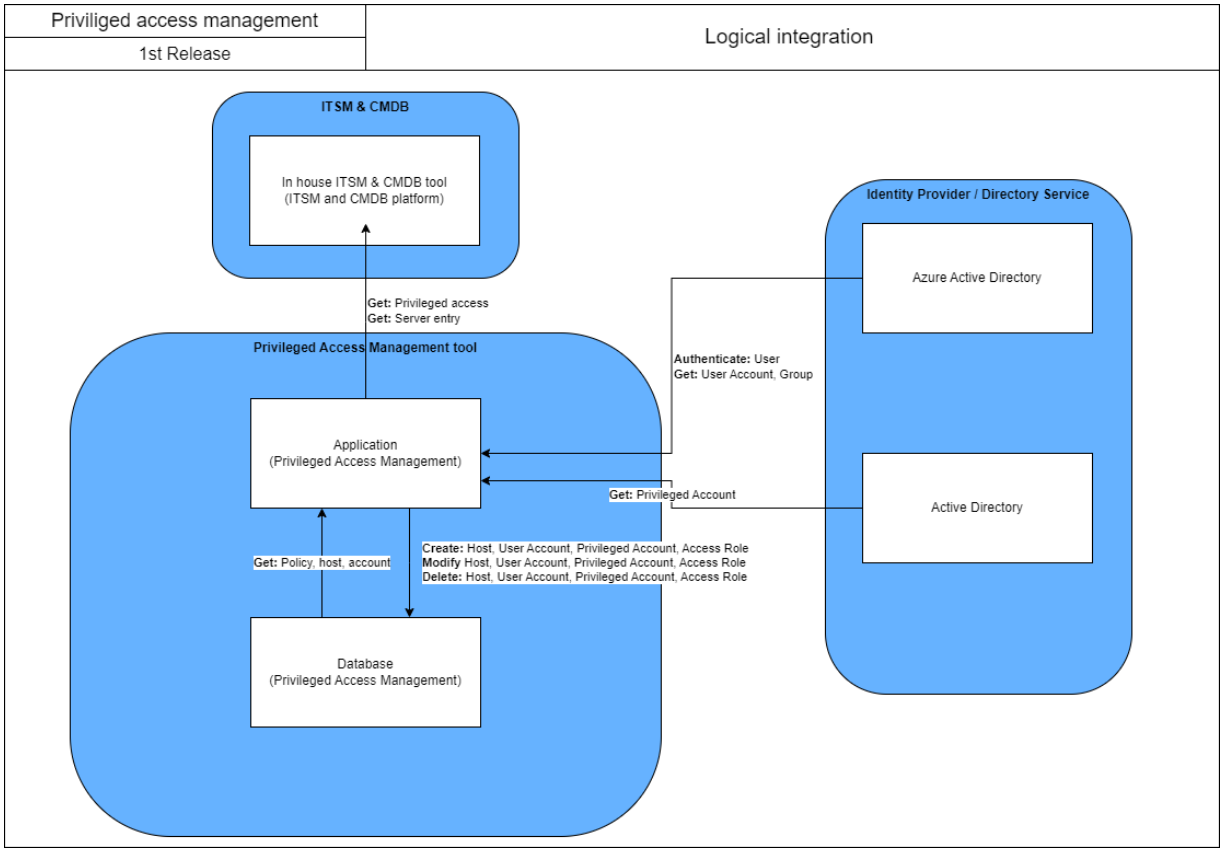

Figure 6, PAM tool dataflow

4 High Availability Design

4.1 Availability

Information technology has had a revolutionary exponential growth in recent decades and with it society and majority of businesses are highly dependent on various IT systems. IT outages will have either direct costs or an estimated costs for organizations, which are dependent on affected IT systems and their criticalities to the organization. [44] The direct costs are associated with repairing possible lost IT systems to continue business operations, or direct penalties in from incidents or unavailability of services. Additional work hours are indirect costs associated with IT outages. IT staff will have to use working hours to solve unexpected problems with IT infrastructure, which reduces their productivity as enabling functions for other aspects of the business to work. Other business is affected as well as some other staff will need IT systems to do their work, creating more downtime for them and reducing efficiency. If thousands of employees cannot perform their duties because of unavailability of IT systems, collectively thousands of work hours have been lost, which has a direct cost to the business itself. The target company is a critical infrastructure provider defined by the NIS2 directive and it provides services, which a modern society requires. Therefore, undisturbed services are not only matter for organization's profitability and reputation. The organization has responsibility towards society to ensure continuity of critical operations and services. The availability of services in modern enterprises generally relies on the methods outlined by Information Technology Infrastructure Library-framework. [45] Involving building a strategy to create highly available IT services to enable the business.

To have high availability, continuity management program is needed. Continuity planning is the creation of systems and processes to ensure that all areas of enterprise is able to perform essential operations or be able to resume the as soon as possible in the event of an incident. Business continuity planning is the overall design of the systems to keep functioning at normal times. IT continuity planning usually sets outlines for criticality of systems, and events, with the focus on services and functionality. This ensures that the system in case has daily high availability and plans for special events are in place. With IT continuity planning comes disaster recovery planning, which focuses on recovering and protecting against disasters. Business availability planning creates mandates for IT continuity, which manages the disaster recovery plans and scenarios.

Importance of an IT service is directly related to its relevance and criticality for an organization. [44] Therefore, it is possible to conclude that privileged access management service is particularly important as it governs accesses to other critical IT infrastructural systems within the organization, inheriting a level of high availability goals from them. This should mean that there is little room for outages, as it would create a cascading effect on to other systems that use the PAM tool. The target company has set an objective for such programs to be available 99.99% of their operation time, meaning the service ought to be unavailable only for 5 minutes monthly. Following this availability objective it is critical that the PAM tool is fault protected. Fault protection ensures that errors or component failures are not noticed by end users or lead into service outages.

Incident is an interruption or a reduction in quality of a service. [46] Criticality of incidents related to the PAM tool are dependent on impacts from the incident, and associated continuity measures (strategies, procedures etc.). If there are no alternative options to bypass the PAM tool, an incident is likely critical. There are usually predefined categories associated with estimated impacts of an incident to support efficient incident management process. Major incidents are incidents that occur to highly critical services, where incidents cause serious interruptions of general business activities and must be resolved immediately.

4.2 Incidents and Disasters

Major incident handling starts as any incident with the it being reported to the service desk or identified during event management monitoring. Regular incidents will require further elevation to a major incident depending on the severity. The incident must fulfil criteria for major incidents defined in the (major) incident management. A major incident manager or team evaluate the incident based on the defined criteria. When the incident is identified, categorized, and prioritized the process to solving the incident can begin. During major incidents that are severe they cause outages, the major incident may include disaster recovery activities if services (tai systems) require rebuilding.

These incidents are not limited to simple outages, and can be cyber incidents with severe safety impacts. Cyber incidents are usually categorized as type of incidents in the incident and major incident management. Cyber incidents have unique characteristics and set of customized processes can be defined into the incident and major incident management to be used in case of cyber incident. Cyber incidents usually occur in real time, and therefore need active monitoring and immediate response by operators and incident managers. Major

incidents and cyber incidents follow an incident response plan, which is a proactive plan to respond to security or other incidents. [47]

Disasters are sudden unplanned calamitous events causing great damage or loss that creates an inability for an organization to serve their purpose for a time period. [48] IT disasters are unexpected major outages that render organizations computer systems or a critical service unusable. Disasters always so severe that basic fault tolerance and high-availability means are inefficient for recovering the system. [44] Disasters are typically physical events that cause an interruption in the service, but can also be major cybersecurity breaches. The key difference between major incidents and disasters is when the management uses disaster recovery plan together with normal incident response plan or major incident management plan. Typically this involves a move from a primary to an alternate location of production. With software services this could mean a different geolocation by several kilometres of physical distance in a different data-centre. In the target organization major cybersecurity breaches are defined as a separate event from disasters, but will follow similar procedures and disaster recovery plan if needed to restore functionality.

4.3 Disaster Recovery Planning

Disaster recovery is meant to serve as a mechanism to dampen the effects of a disastrous event on a service, ideally restoring its functionality to the original level. IT disaster recovery planning is typically a part of risk management and continuity planning as majority of organizations require IT systems. [49] Furthermore, most of the cybersecurity policy frameworks used at the target organization demand, or highly recommend using a disaster recovery plan to restore systems that have been impacted by a disaster. Therefore, it is essential that a plan for disaster recovery is created. In a case of emergency restoring systems will be more difficult during disasters without comprehensive disaster recovery plans.

An initial part of continuity management and disaster recovery planning is to perform business impact analyses. Generally an approach and methodology is defined and tailored for an organization to identify and assess (e.g., business impact analysis) potential impacts from different type of disasters. Business impact analyses will allow identification of areas and business functions impacted by downtime of a service Planning of rational and efficient disaster recovery including creation of a plan can be initiated based on information from business impact analyses. IT disaster recovery planning should follow and support overall business continuity planning. [50]

Next steps in creating a disaster recovery plan is to assess the likelihood and potential consequences of the risks associated with the service. Continuity risks have been physical events in history, for example earthquakes, floods, fires etc. Recently, the relevance of IT continuity has risen, which is for example outlined in the chapter 3 for risks associated with privileged access management. The risks should be assessed for a variety of scenarios in addition to natural disasters, simple equipment failure is a big risk, sabotage, insider threats are relevant. In addition to physical risks, variety of other continuity risks should identified and assessed. These risks are for example equipment failures, sabotage and insider threat. The impacts from a risk can be calculated by factoring in the direct financial losses due to faulty revenue generating activities, the brand's reputational damage, employee productivity, health and safety, and progress towards business initiatives or goals. By factoring in all these, it enables rationale and effective planning of disaster recovery measures and procedures based on the associated impacts.

The following part in creating the disaster recovery plan is to document all dependencies of the service. This is done by creating a complete inventory of hardware and software, and other assets related to the service. If one of these parts is unavailable, will the whole service be unavailable? With the chosen privileged access management tool, one of the aims of the application was to reduce dependencies in the target systems, meaning the PAM tool can work independently. The tool is mainly dependent on common (capacity, network, computing, storage etc.) and infrastructure services which is creating a dependency on this software. The application instances are hosted with virtual machines, that are reached via network services (DNS, routers, firewalls etc.). This creates dependencies on these systems to work to reach the PAM tool. In case of a disaster, it is possible that these dependencies can be circumvented with the use of break glass processes and procedures. Break glass is a term in IT used to describe solving of a catastrophic event, by destroying a metaphorical glass of a fire alarm. An option for privileged access management is to have break glass accounts and credentials, which can be retrieved in case of disaster or outage of the PAM tool. [51]

In addition to the above, Recovery Point Objective (RPO), Recovery Time Objective (RTO) and Recovery Consistency Objective (RCO) targets need to established for IT continuity based on information from business impact analyses. Recovery time objective is the maximum amount of time taken to restore the service into function after a service disruption. Calculating the RTO starts from the point of disaster. The point in time where the full recovery of lost systems is reached must be measured to have reasonable conclusions

from the RTO. The recovery point objective is the maximum time window where the information or data could have been lost. In order for the service to resume regular operations the RPO must be reached. There are variety of backup strategies and frequencies, which are partially determined by RPO targets. RCO is a target metric that indicates the amount of inconsistencies in recovered data, for example missing entries or corrupted entries, which can be tolerated by an organization. [50] The target company is not using RCO targets in their continuity program, but is generally an used metric.

4.3.1 Service Resilience

If disaster recovery planning is partially relying on resilient of an IT environment and service, robustness and redundancy must be incorporated into design of the service and the supporting infrastructure. Three main approaches for resilient and redundant system architecture are a shared system, hot standby, and cold standby system. [44] As physical disasters or crisis situations may occur outage types, a varying degree of geo-local separation must be applied to the architecture to avoid multiple different loss scenarios depending on the location.

The shared system approach to architecture states, which the service ought to rely on multiple independent subsystems that provide the same service for the end user. [44] This way, the services are kept available for users through different channel, if one of the subsystem has become unavailable, since they are not impacted by even a complete loss of one subsystem. In addition, this enables that the service should be separated into many different instances on multiple sites. Additional sites in return would raise the cost of the implementation and therefore must be acknowledged, when creating a budget. Another problem with this approach is that the application servers must be active at all times, even when very little or no traffic at all would be outgoing. The maintenance work that this approach generate would be much more, as all the systems need to be synchronized to apply configuration updates, and considering that they must be geologically differentiated, clustering the servers would be out of the question. This leads to tons of manual work for administrators of this service and possibly negatively affecting the availability goal of the service. [44]

The other two approaches are hot-hot standby and hot-cold standby setups. The primary system is used in both approaches for users if it operates normally. If the primary system is not operating as designed or it is unavailable, the secondary system will be used. [44] These standby systems can be categorized into “hot” or “cold” or sometimes “warm” depending on

their behaviour. A hot system is identical and operational at all times with the primary system, and a cold system is not running and is only brought up in the case of an actual disaster. The distinction between cold systems and warm systems is that a cold system is constantly offline and a disaster recovery process must be started from installing required software on to the system, but a warm system is somewhat prepared for a disaster by frequent updates on the required software. This distinction is disputed in some sources, but used in the target organization.

4.3.2 Cost Versus Benefit

It might seem trivial to ensure business continuity through a disaster recovery plan, but the challenge with preparing for worst cases is to balance the cost to benefit. In principle, costs of disaster recovery mechanisms and procedures are balanced with associated impacts in a mature continuity management program. “ A design for disaster recovery is a compromise between financial expenditure and the redundancy achieved”. (Schmidt, 2006) [44] As disaster recovery scenarios are unlikely, preparing for them might seem expensive. Therefore, a balance between allocating resources for recovery measures and the benefit that the disaster recovery process gains has to be considered. With unlimited funds, previously mentioned RPO and RTO targets are irrelevant as there is no incentive to not have everything duplicated at real time. The initial cost evaluations of business continuity of the tool are confidential and made by the risk assessment group in the target organization therefore, they cannot be stated outright in this thesis.

5 Disaster Recovery Scenarios and Implementation Plans

Key takeaways from the knowledge base of the previous chapters are as follows. **Identity and access management** is used to manage and control access to assets and resources within an organisation's IT environment. **Privileged Access Management** is a major feature in IAM is to manage privileged accesses, that are specially vulnerable to **threats and cyberattacks** as they provide adversaries unlimited access for lateral movement. The PAM solution should complete the requirements of **frameworks and legislations** to improve the cyber resilience of an organization. The solution should be accessed through **digital identities** that represent an individual user, where the identity is **authenticated** with **multifactor authentication**, including **credentials** and **biometric authentication**. These identities are then **authorized** through a hybrid between **role-based access control** and **attribute-based access control** to access the target resources. The then accessible privileged access management solution must be made **highly available** through **service redundancy** and resilient to incidents and **disasters** to ensure **IT continuity**, while being reasonably **cost-efficient**.

The way of working during the implementation project at the target organization was to approach high availability and disaster recovery from the point of view of disaster scenarios. Part of the design science research artifact, which is the creation a service recovery plan for the privileged access management tool in the target organization is to define and estimate some disaster recovery scenarios. These scenarios are events that have the potential to make the service unavailable and obligated for a disaster recovery procedure additionally to normal functions. The scenarios gather information from the knowledge base leveraging its information for the decision making process within the implementation project.

This chapter will provide the information that lead to the creation of the research artifact.

5.1.1 Component Malfunction or Unavailability

Malfunction or unavailability of a critical system component will result that the system or its key functions are not available. The key principle in disaster recovery planning is to increase resilience with redundancies if high-availability is a requirement for the system. Redundancy can be achieved through providing a backup components. It is a basic, but very reliable way to have high availability for the service. The issue with having backup components is that the costs are linear, every backup component has a similar cost in implementation as the original

one. It is possible to have multiple backup components, and therefore the term duplication is not necessarily the most accurate.

The goal for resilient system architecture and redundancies is to reduce single point of failures. Systems have usually components or dependencies that create single point of failures (SPOFs). A system or its key features will be unavailable if there is an error in a single point of failure. The identification of single point of failures requires mapping of critical assets associated with the system. The critical assets are for example physical data centers, network cabling, physical servers, logical infrastructure services and servers. The assessment and conclusions on potential single point of failures are dependent from responsibilities and scope of the disaster recovery planning. In Figure 7, the all the critical components and their planned redundancies are illustrated, excluding the user and the target resources. The application server instances and database instances were identified as critical components, and therefore there components are at least duplicated. The web access components are not critical and the duplication is used to ensure the performance. Only non-redundant component was found to be the loadbalancer, which diverts the connections to the application instances depending on the connection loads and is not a direct component of the PAM architecture, instead being a component of the network in the organization. Additionally, the loadbalancer is not duplicated as users may connect the application server instances directly

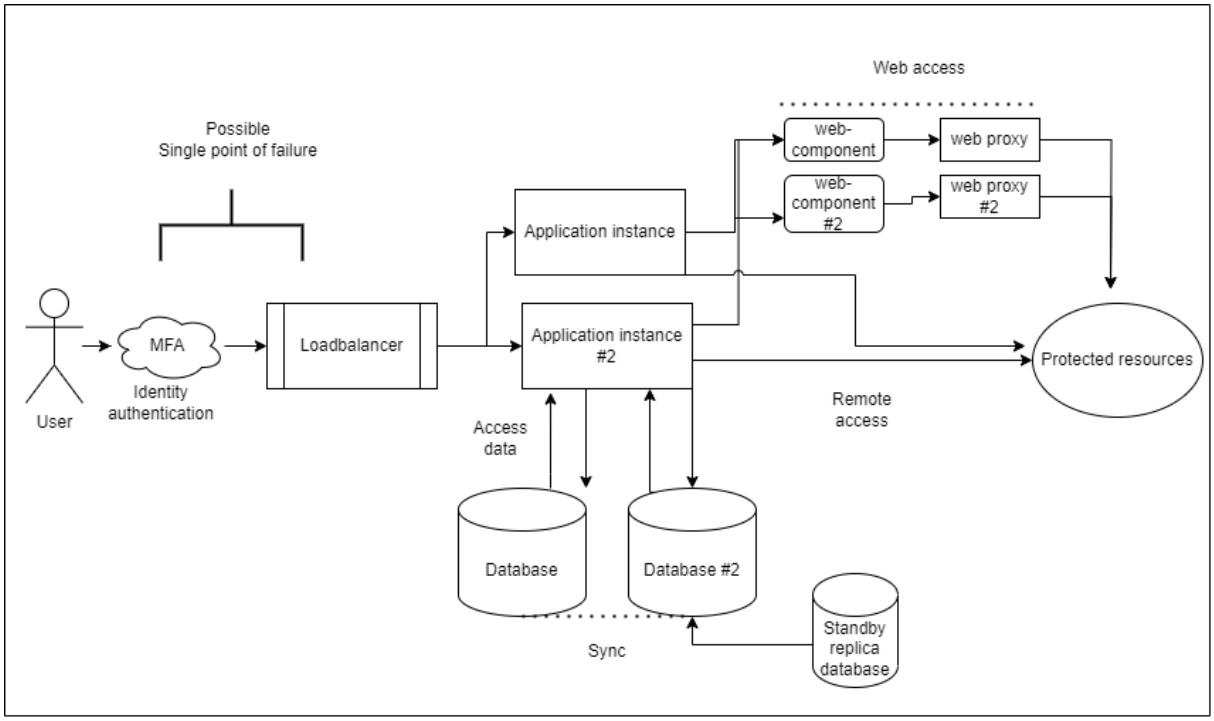

Figure 7, Redundant PAM architecture with SPOF's

Reliance on multifactor authentication by an external third party, creates a possible single point of failure, where the application cannot be reached if the service is unavailable. The multifactor authentication services are leased as software-as-a-service and therefore is not possible for the project team to make highly available. This is not expected to be and issues as the provider has contractual availability and reliability requirements. It has to be considered when designing and implementing a resilient architecture and planning disaster recovery procedures. One possible way to avoid the single point of failure, is to circumvent it through a rigorous break glass procedure that uses local authentication if the service is unavailable.

The application instances are hosted in virtual machines that are duplicated to be hot & hot, with session persistence by identity. If a connection is made by an identity towards the load balancer, the load balancer chooses the primary application instance and creates a persistent session for the identity towards the application instance. Persistence demands that the session will not swap between application instances while it is active, avoiding possible synchronization errors with the application state. This enables the session from the load balancer to be only on the primary application instance, which makes it possible to not have synchronization between the application instance one and application instance two, reducing the work needed by the application instances. The lack of synchronization will mean that the sessions are not possible to be transferred while the application is running in a case of unavailability on the primary application instance, and a new connection has to be initialized. During the implementation project, the lack of synchronization was considered to be an issue with user experience, however it was later deemed as a necessary feature of the chosen service could be made scalable as more application instances could be added anytime.

The application instances forward the connection towards the protected resources when they are, RDP or SSH remote-connection, if they are HTTPS-based web accesses the connection will go towards the PAM tool's web-components and forward to web-proxies. The web-components create a secure session using the access data from application instances to the web-proxies, which allow the users to use a blank web-browser from the proxy to access the protected resources through a web-application. One carrier is used as a primary one, while the other is a backup secondary one, to increase resilience. Similarly the web proxies are duplicated to have a primary and a secondary one for fault tolerance. With this fault tolerance, the application instance will be configured in a way that it redirects the connection towards the secondary component if the primary one is unavailable, similarly to the load balancer.

The database environment consist two database instances and the instances are configured that the both instances are active. Another of the two database instances will be configured as the primary node, which will support the production environment. The secondary database instance will act as supporting node and it will be set as the primary instance based on heartbeat from the database environment. This ensure that the system and services will be available if there is a failure with the another database instance. One read-only cold standby replica was deemed necessary to ensure that a complete data loss will become extremely unlikely, as the standby replica will only be synchronized rarely to store longer backups of the databases, that can then be used to restore the original state of the application in case of a disaster.

To achieve high availability through the redundant components a load balancer is needed for the system. As stated previously, the load balancer will efficiently distribute the connections towards the application instances automatically. [52] The load balancer ensures that in a case where one application instance is offline, the connections are re-directed to an online application instance, preventing outages. The load balancer is not considered a component of the PAM tool, as it is an IT infrastructure for other applications in the organization also. However, it is a critical components when considering the high availability of the service and it needs to be included in the design documentation. Load balancer uses different types of load balancing algorithms: [53]

- Round Robin – Requests are sequentially distributed between the application instances
- Least Connections – New requests are sent to the application instance with the fewest current connections to the clients.
- Least time – Requests are sent to whichever application instance has the fastest response time
- Hash – Distribution based on a defined secret
- Random with two choices – Picking a server randomly

It was concluded that generally Round Robin routing algorithm should be used as the server and port information is preserved for the session duration. For the high availability design to work it needs to reach a pre-defined URL and regex match response with an expected string to determine the availability of the underlying application instances.

5.1.2 Data Loss and Backup Failure

Data loss can be prevented by designing backup strategy aligned with RPO, RTO and RCO targets. The backup strategy is implemented with the backup setup and configuring the backup services based on the backup strategy, where data can be recovered even if it is considered lost, corrupted, compromised or stolen through hardware failure, human error or malware. The plan is to identify the essential data for the PAM tool to keep operational and to ensure it can be accessed even in a disaster. Duplication of the PAM tool database is a part of the approach to ensure resilience if the another database server instance is lost. The primary database server is asynchronously replicated to the standby database server.

Data loss can occur if the database becomes unavailable or corrupted, but data loss is not limited to component malfunction. Data loss can occur from a successful ransomware or other cyber-attack on the database of the service The service needs to be air gapped on the network level to ensure that the breach is contained in case of a breach, which means that the service can be separated in the network level from the outside internet as well as the organizations own network. If a corruption or data loss is detected in the primary database, the heartbeat system should not synchronize it with the secondary one and a backup can be driven from the secondary system to restore the primary one. However, it is a possibility that both the primary and the secondary one become unavailable or corrupted. If it were to happen, there is a standby read-only database to provide a backup to restore the state before the corruption. The standby read-only cannot be used in case of an outage as a replacement for the primary or the secondary one, but the primary or the secondary one can be restored and rebuilt from the standby one.

The databases will be implemented into two separate datacentres. In the data centres the databases will be under the same cluster, where the primary data base will be readable and writeable, and the secondary one similarly, but only the primary one is writeable from the user perspective. The standby database one is read-only, it removes the possibility of both databases becoming corrupted, as the standby one cannot be written over. The secondary one synchronizes with another standby one only from time to time to keep the backup data up to date. In figure 8, the architecture is illustrated to highlight the scalability of the setup, by having a possibility to add more databases or application instances if required.

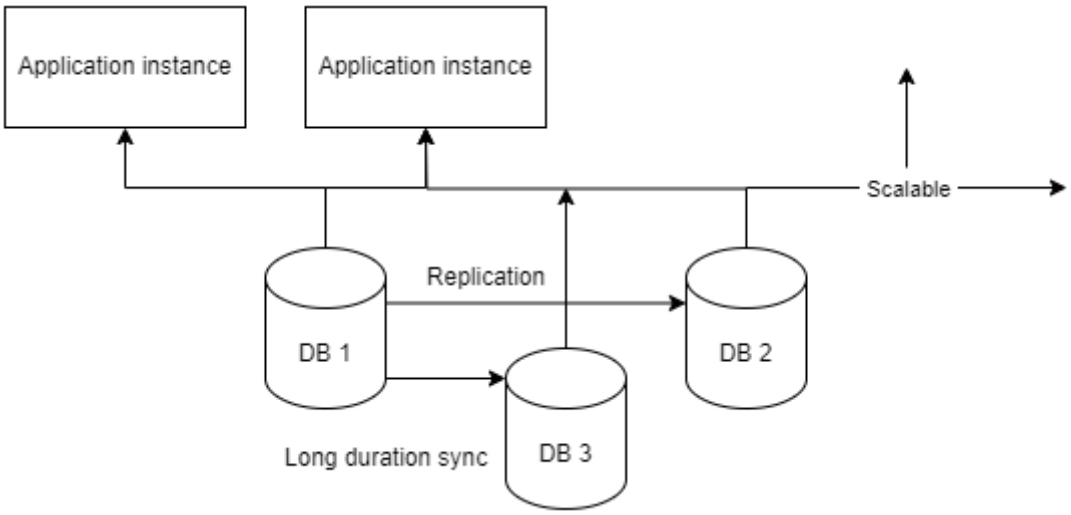

Figure 8, Database set up

Originally the database replication was suggested to be even more redundant, but it will have to be cut due to concerns it being not reasonably budgeted. The primary thought for conserving resources, was that the break glass procedure will reduce the risks involved with data loss and additional databases will reach the law of diminishing returns.

5.1.3 Stolen Credentials or Certificates

A cybersecurity focused disaster that was considered during the project, was that the main access credentials or certificates of the PAM tool may get stolen. If stolen, the whole system would be inaccessible and in danger. The original plan is to monitor credentials and certificates related to the access to the PAM tool in a Security and Information Event Manager (SIEM). SIEM would then be able create an alert, that would automatically halt access through the credentials or certificates to the system. This was based on a regular incident response to a cybersecurity breach, where the main goal is to isolate and contain. Identifying, revoking, communicating, and forensic analysis would be left to the Security Operations Centre (SOC), since they are responsible and capable of handling similar events in the organization. However this scenario is out of scope, as preparing for it is dependent on the progress of the project and it will be only possible to consider after the system is made highly available.

5.1.4 Personnel Disaster

Personnel disasters are more part of the business continuity planning at the target organization, however they are important to consider for disaster recovery. The recent Covid-19 pandemic highlighted the need for preparedness against workforce disasters. [54] One important scenario that will eventually arise is that the workforce might be unavailable due to a disease for a long period, or all of the personnel with the knowledge of managing the product are unavailable due to some unspecified reason. For this situation it is useful to have redundancy within the workforce, in a way that the service can stay available even if many of the personnel become unavailable. This can be done using automated services and having remote access to all of the critical components in the system and training multiple personnel to manage the system. The PAM tool will be operated entirely remotely if using the organizations own network, so there is already built in redundancy towards personnel unavailability. The need to make it possible to access from outside the network will be out of scope for the test environment of the service, but might be useful to note when considering the production environment. Proper documentation will provide the possibility to avoid knowledge loss if a worker needs to be replaced or retrained in case of a workforce disaster.

5.1.5 Vendor Disasters

The vendors that provide services for the solution need to have their own disaster recovery in a case of disaster at their facility. It is possible that they may have a problem with for example electricity outages or security breaches, however it cannot be in our services scope to prepare for them. Some assumptions on the level of service continuity provided by the hosting have to made:

Security has to be taken into consideration by the vendors as they are liable to secure physical access to our services. For example they should be prepared for fire emergencies, natural disasters, or sabotage by an outside threat. Therefore, they must create a comprehensive plan to secure the resources that are in their responsibility.

Electricity outage resilience can be achieved by the vendors using back-up power generators. These stand by generators provide an alternative electrical load to the servers by automatically detecting power outages and turning on in a case of outage. Most commonly back-up power units run on diesel and generate enough power to keep the service running during the power outage.

Vulnerabilities in the application are possible and the service vendor is liable to keep the application up to date to avoid vulnerabilities in the software. An attacker could leverage a vulnerability in the software application to gain unauthorized access or manipulate configurations to access sensitive data. Therefore, the vendor should establish a robust vulnerability management program to identify and remediate vulnerabilities in their provided service.

5.2 Break Glass Procedure

Break glass procedures are relevant for the disaster scenario, where the entire provided service is unavailable and users have a need to access protected resources within the application. The protection and restriction associated with the PAM tool must be able to be bypassed in case the PAM tool is not available or able to connect to the protected targets. The PAM tool will not have cascading effects on the protected assets, when sufficient break glass mechanism and procedures are defined, implemented and tested. The break glass enables accessing of the protected assets for example with the disaster recovery accounts or credentials. The following procedure is based on the target company's own best practises for a break glass and incident response loosely adopted from NIST. [55]

Creation and provision of break glass accounts are based on pre-defined procedures or accounts are already created into the protected assets. Access to break glass accounts potentially need additional procedures, which can be for example firewall openings or similar type of actions. If impacts from a disaster are significant, there is a need to bypass the privileged access management process through the PAM tools. Use of break glass procedures generally require accessing the protected assets with accounts that are provisioned into the protected targets beforehand. This requires that authentication details (passwords etc.) associated with the accounts are available. There are many approaches to share and store authentication details of the break glass accounts. The authentication details of disaster recovery accounts used in the break glass procedures are distributed that those are available in case of disaster. The authentication details can be distributed and stored in disaster recovery plans, which are stored in a physical safes. In addition, the authentication details can be stored in virtual vaults if those are sufficiently resilient. [51]

The principle for using break glass procedures and accounts is ongoing incident during the protected assets are not available through the PAM tools. There are approaches for limiting usage of break glass procedures and accounts. The most common approaches are

limiting use of break glass procedures by limiting access to the protected assets in other layers, for example with firewall rules. Additionally, access can be limited by restricting access to authentication details in combination with monitoring use of the break glass procedures and accounts. The initial monitoring is usually performed by a SOC services. Authentication details of all accounts including break glass accounts are generally sent to a security information and event monitoring tool, which support the SOC services. The SOC services have defined procedures, which they will execute if a break glass account is used. The procedures may include closing of connections if use of the break glass account cannot be traced to any ongoing incident management activities. If use of a break glass account cannot be logged at protected asset and monitored in real-time, the monitoring can be performed from a secondary source. The secondary source for monitoring can be access log for a physical safe or digital secret vault. Each use of these emergency accounts will be reviewed and that use of account is associated with a valid reason and conclusion on appropriateness of the usage. The review results are documented as an evidence.

The post incident activities may include changing authentication details of break glass accounts, disabling of break glass accounts, or deletion and re-creation of break glass accounts, which prevents use of the accounts after the incident. Additionally, the post incident activities can include review of performed activities and analysis on effectiveness of the performed disaster recovery activities. Results from the analysis are reflected to the disaster recovery plans.

It is important to note that break glass procedures are a powerful tool and should be used judiciously. There must be a balance to the need for emergency access with the imperative to maintain a secure and auditable environment. The implementation of break glass scenarios should align with industry best practices, framework recommendations, regulatory requirements, and the organization's overall security strategy.

6 Evaluation

The evaluation of disaster recovery scenarios, which are the research artifact of the thesis need to be limited to testing the high availability of the application instances in the privileged access management tool. This thesis was a part of an ongoing project in the target company, which created a scheduling conflict with the planned thesis completion date and the progress of the ongoing project. Due to this conflict, only the application instances could be made redundant and ready to be evaluated. The original plan was to test and evaluate the resilience of all the components in the PAM architecture.

6.1 Application Instance Shutdown

Evaluation of the most important aspect of the PAM service availability is to test whether an application instance is available even if one of the instances suffers a disastrous incident. The evaluation was done over a period of a week with the IAM team evaluating the results of a controlled disaster scenario. The scenario consisted of a situation where the primary application instance was periodically shut down, to test whether the load balancer was able to transfer the connection load towards the second instance. Similarly to test the redundancy of the secondary fallback application instance. Figure 9, shows the starting status of the application instance and figure 10 shows the end status of the application instances during the test.

The disaster recovery plan for this scenario consists of instant RTO and no data loss RPO. The recovery time should be instant as the fall back system is not working as intended if there is a delay to access the service. The RTO of an application instance being rebuilt in case of an outage could not be tested, as the test was to only assess the application instance availability. Recovery point objective is that no data should be lost, as the application instance should not be handling any data that can be lost. The hypothesis is that as the primary

application instance goes offline, the secondary instance is immediately usable, but the user will have to reauthenticate, as the application instances are independent of each other.

Figure 9, starting the disaster scenario

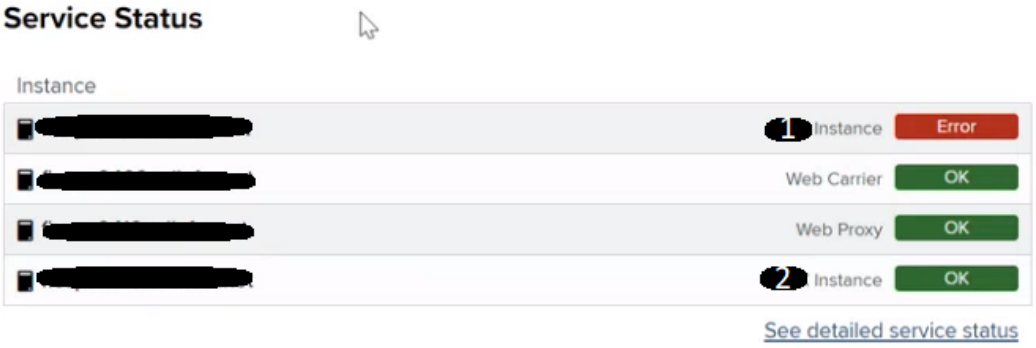

Figure 10, ongoing disaster scenario

Immediate findings from the service becoming unavailable were, that the load balancer was able to redirect new connections to the secondary application instance immediately, but active connections were not able to switch instances. The hypothesis was that the load balancer was incorrectly configured to poll the active connections only once a minute. The hypothesis was revealed to be incorrect during an interview with the technician that configured it, as the load balancer was set to poll the availability of the instances 84 times a second. This led to another hypothesis that the web browser, which was used to access the application instance was honouring the cookies sent by the original connection. This hypothesis was deemed to be the reason for the incorrect operation of the application by the technician. The incorrect operation

was able to be fixed, however this was only an user experience issue, as the underlying secondary component was assessed to be fully functional.

Implementation of the failover system effectively means that only the connection with the target servers will be shut down, but the session within the server will be able to continue operating once the new connection is routed through the secondary system. Therefore, no data loss is possible to occur in the target resources during an outage of the application instances. The application instances only handle access data, which will be examined individually every time a new connection is made.

The implications of correct operation of the failover system is that the environment is able to withstand an outage in one of the application instances. However, it is possible for both of the components to fail simultaneously, which would mean that the whole PAM solution would need to be circumvented. This will be achieved through a break glass procedure, where an administrator would access the target servers directly, until the PAM solution can be rebuilt.

6.2 Database Unavailability

Testing the database availability could not be done, as they could not be implemented in time for the completion of this thesis. However, we can hypothesize how the redundancy would have benefitted the service. The two databases were designed to be in an active-active configuration where they would run simultaneously, but only one of them being the primary one. In this configuration, if one database becomes unavailable due to failure or disaster, the other can seamlessly take over, minimizing downtime and ensuring continuous service availability. If a disaster would occur in the primary database, the process would typically involve promoting the secondary database to a primary one, reconfiguring the system, updating DNS records, and redirecting the application instances traffic towards the new primary database.

The benefit of this redundancy is that the service avoids a single point of failure in the access data, providing possibility for disaster recovery for the service. The standby backup database would provide an additional layer of redundancy further safeguarding against data loss and ensuring quicker recovery times. Therefore, the RTO would be reduced and RPO should be always possible to set before the disaster occurs.

6.3 Break Glass Scenario

Break glass scenario is typically used in situations where standard access to critical systems or data is not possible. Accessing the critical resources should always proceed through the PAM application, however it is possible that the application is unavailable due to disaster and accessing the resources is impossible. If any instance of the application cannot be reached, must the target resources be available to use. The workflow of the break glass scenario will be done by the security operation centre's giving guidance to an incident response team.

In a hypothetical break glass scenario where all application instances are unavailable. There will be pre-staged credentials and monitored break glass emergency credentials for the administrators to use. These pre-staged credentials will only be accessible through incident response team, who will authorize them by their judgement on the severity of the incident. A dedicated mechanism to trigger a break glass scenario will not be implemented into the application itself, instead the trigger will always have to be done manually, due to the nature of the scenario. While being manual work, some level of automation can be applied, since the process can utilize virtual vaults and safes for storing the credentials.

Every action taken with the break glass credentials will be monitored and logged, since after the incident is resolved, there will be a review process for all the actions taken. This monitoring and reviewing work will be done by the incident response team. To end the break glass procedure, the credentials will be deleted indefinitely.

A short test of local authentication to the protected resources was done, however it did not include the future break glass procedure. The short test was to simply reach the protected resources while the authentication service was unavailable. It did prove that the protected resources will be reachable even if the PAM solution or the multifactor authentication service is unavailable. The test was ran by using local authentication credentials on the target resources.

7 Conclusions

Managing privileged access is essential part of identity and access management and to manage privileged accesses many compounding structures need to be in place for a continuous delivery of the service. To enable resilience for a system that manages privileged access, it need to be made highly available to ensure that no cumulative interruptions to other services that enable efficient energy production in the target company.

The research questions set for this paper were:

- What are privileged access rights?
- Why should privileged access rights be secured?
- What type of protection mechanisms can and should be applied to them?
- How the resilience of these protection mechanisms is ensured in case of a disaster?

The answers to them can be summarized as:

Privileged access rights are limited access rights to resources that are deemed critical or authoritative in nature, where an elevation of privileges is required compared to regular access rights to perform administrative tasks or access restricted resources.

Securing privileged access rights is a fundamental component of a comprehensive cybersecurity and identity and access management strategy. They ought to be secured as it helps to protect sensitive information, prevents unauthorized actions and ensures overall resilience towards cyberthreats.

The way to effectively protect privileged access rights against threats is through privileged access management. PAM solutions are designed to protect critical systems, sensitive data, and infrastructure by enforcing strict access controls, monitoring privileged activities and implementing security best practises through frameworks.

Resilience of the PAM solution can be achieved through creating disaster recovery-ready solution, that should achieve >99% uptime. Disaster recovery-ready solution is achieved through fault protection, redundancy, and IT continuity planning through a disaster recovery plan.

The thesis followed the set parameters and leveraged information from a knowledge base, adopted from design science research and produced a research artifact for the target company by providing documentation for high availability and disaster recovery. The design science research proposed two questions: “What utility does the artifact provide?” and “What demonstrates that utility?” These questions are answered by the service recovery document provided for the target company, that states the ways the service has been made highly available and recoverable. Demonstration of that utility is that the documentation is required by the organization’s compliancy requirements towards the legislation in the European union.

7.1 Limitations and Future Study

This thesis is limited to one particular case of implementing a privileged access management solution in a target organization, which might not be applicable in other organizations or with a different solution choice. By working with a target organization, many of the details had to be anonymized making the thesis less specific. However, the generality of the thesis provides an excellent opportunity to see how and why a similar tool can be made resilient.

As the underlying conditions in the project at the target organization brought up challenges with scheduling with this thesis, the evaluation could not be made for the resilience of the planned database redundancy. Other component’s availability are set to be tested during the following year.

The research on the possible business impacts of the risks associated with privileged access management could not be used in the thesis. The quantifiable facts and logic would reveal sensitive information of the target company, but would have been useful for further evaluation of the reduced business risk by the PAM solution.

This research could be extended to measure all of the disaster scenarios proposed in the implementation chapter and possible further scenarios, with their relevant business impact analysis. This would create a more comprehensive look at the possible disasters and challenges that may occur while implementing high availability for a privileged access management tool.

References

[1] A. C. Keung. Ng, *Contemporary identity and access management architectures : emerging research and opportunities*. Hershey, Pennsylvania (701 E. Chocolate Avenue, Hershey, Pennsylvania, 17033, USA): IGI Global, 2017.

[2] J. vom Brocke, A. Havner, and A. Maedche, *Design science research. Cases*. Cham, Switzerland: Springer, 2020.

[3] A. Hevner *et al.*, 'Design Science in Information Systems Research', *Manag. Inf. Syst. Q.*, vol. 28, p. 75, Mar. 2004.

[4] E. Bertino and Kenji. Takahashi, *Identity Management : concepts, technologies, and systems*. in Information security and privacy series. Boston, Massachusetts ; Artech House, 2010.

[5] 'Microsoft Digital Defense Report 2022', 2022. [Online]. Available: https://www.microsoft.com/en-us/security/business/microsoft-digital-defense-report-2022

[6] 'The 2023 Global Ransomware report', Fortinet, 2023. [Online]. Available: https://www.fortinet.com/content/dam/fortinet/assets/reports/report-2023-ransomware-global-research.pdf

[7] S. Murugesan and I. Bojanova, *Encyclopedia of cloud computing*. in Wiley - IEEE. Chichester, West Sussex, United Kingdom Hoboken, NJ: John Wiley & Sons, 2015.

[8] 'DIRECTIVE (EU) 2022/2555 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 14 December 2022'. Accessed: Oct. 10, 2023. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32022L2555&from=EN

[9] O. Rantala and T. Kievari, *Information Security Strategy for Finland The World's Most Trusted Digital Business Environment*. Liikenne- ja viestintäministeriö, 2016. [Online]. Available: https://julkaisut.valtioneuvosto.fi/handle/10024/75353

[10] International Organization of Standardization, *ISO/IEC 27001:2022*, 3rd ed. 2022. [Online]. Available: https://www.iso.org/standard/27001

[11] *NIST SP 800-207 Zero trust architecture*. 2020. Accessed: Oct. 31, 2023. [Online]. Available: https://csrc.nist.gov/pubs/sp/800/207/final

[12] S. Teerakanok, T. Uehara, and A. Inomata, 'Migrating to Zero Trust Architecture: Reviews and Challenges', *Secur. Commun. Netw.*, vol. 2021, p. 9947347, May 2021, doi: 10.1155/2021/9947347.

[13] M. Gupta, J. Walp, and R. Sharman, Eds., *Threats, countermeasures, and advances in applied information security*. Hershey, PA: Information Science Reference, 2012.

[14] One Identity, 'What is IGA (Identity Governance & Administration)?' Accessed: May 11, 2023. [Online]. Available: https://www.oneidentity.com/what-is-iga/

[15] NIST, 'Identity assurance level'. Accessed: Sep. 11, 2023. [Online]. Available: https://csrc.nist.gov/glossary/term/identity_assurance_level

[16] *What is considered personal data under the EU GDPR?* Accessed: Sep. 11, 2023. [Online]. Available: https://gdpr.eu/eu-gdpr-personal-data/

[17] *Advances in user authentication*. New York, NY: Springer Berlin Heidelberg, 2017.

[18] M. P. Eve, *Password*. in Object lessons. New York: Bloomsbury, Bloomsbury Academic, An Imprint of Bloomsbury Publishing Inc, 2016.

[19] *CCS '20: proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security : November 9-13, 2020, virtual event, USA*. New York, New York: Association for Computing Machinery, 2020.

[20] S. Mahnken, 'Today's authentication options: the need for adaptive multifactor authentication', *Biom. Technol. Today*, vol. 2014, no. 7, pp. 8–10, 2014, doi: 10.1016/S0969-4765(14)70126-2.

[21] *Recent Advances in Biometrics*. IntechOpen, 2022. doi: 10.5772/intechopen.97986.

[22] Okta, 'Benefits of single sign on'. Accessed: Sep. 20, 2023. [Online]. Available: https://www.okta.com/uk/blog/2022/04/benefits-of-single-sign-on/

[23] A. Stavrou, H. Bos, and G. Portokalidis, Eds., *Research in attacks, intrusions, and defenses: 17th international symposium, RAID 2014, Gothenburg, Sweden, September 17 - 19, 2014 ; proceedings*. in Lecture notes in computer science Security and cryptology, no. 8688. Cham: Springer, 2014.

[24] 'ISO 27002:2022 5.12'. 2022. [Online]. Available: https://www.isms.online/iso-27002/control-5-12-classification-of-information/

[25] T. R. Weil and E. Coyne, 'ABAC and RBAC: Scalable, Flexible, and Auditable Access Management', *IT Prof.*, vol. 15, no. 03, pp. 14–16, May 2013, doi: 10.1109/MITP.2013.37.

[26] V. Hu and V. C. Hu, *Attribute-Based Access Control.* in Artech House information security and privacy series. Boston, Massachusetts: Artech House, 2017.

[27] H. Rasouli, *Proposing a Digital Identity Management Framework: A Mixed-method Approach.*, vol. 33. in Concurrency and computation, vol. 33. 2021.

[28] A. C. K. Ng, *Contemporary identity and access management architectures: emerging research and opportunities*. Hershey, PA: IGI Global, 2018.

[29] I. Indu, P. M. R. Anand, and V. Bhaskar, 'Identity and access management in cloud environment: Mechanisms and challenges', *Eng. Sci. Technol. Int. J.*, vol. 21, no. 4, pp. 574–588, Aug. 2018, doi: 10.1016/j.jestch.2018.05.010.

[30] E. Sindiren and B. Ciylan, 'Application model for privileged account access control system in enterprise networks', *Comput. Secur.*, vol. 83, pp. 52–67, Jun. 2019, doi: 10.1016/j.cose.2019.01.008.

[31] P. K. Manadhata and J. M. Wing, 'An Attack Surface Metric', *IEEE Trans. Softw. Eng.*, vol. 37, no. 3, pp. 371–386, Jun. 2011, doi: 10.1109/TSE.2010.60.

[32] 'COUNCIL DIRECTIVE 2008/114/EC of 8 December 2008 on the identification and designation of European critical infrastructures and the assessment of the need to improve their protection'. Accessed: Aug. 10, 2023. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32008L0114

[33] J. Soldatos, *Cyber-physical threat intelligence for critical infrastructures security : a guide to integrated cyber-physical protection of modern critical infrastructures*. in Cyber-Physical Threat Intelligence for Critical Infrastructures Security. Hanover, Massachusetts: Now Publishers, 2020.

[34] 'Kansallinen riskiarvio 2023 Sisäinen turvallisuus | Sisäministeriön julkaisuja 2023:4', Valtioneuvosto, Sisäministeriö Helsinki, 2023. Accessed: Aug. 23, 2023. [Online]. Available: https://julkaisut.valtioneuvosto.fi/bitstream/handle/10024/164627/SM_2023_4.pdf?sequence=1&isAllowed=y

[35] 'Energy prices in Eu 2023'. Accessed: Sep. 30, 2023. [Online]. Available: https://www.consilium.europa.eu/en/infographics/energy-prices-2021/

[36] S. Zilincik and I. Duyvesteyn, 'Strategic studies and cyber warfare', *J. Strateg. Stud.*, vol. 46, no. 4, pp. 836–857, 2023, doi: 10.1080/01402390.2023.2174106.

[37] Crowdstrike, 'Lateral movement', Cybersecurity 101. Accessed: Dec. 09, 2023. [Online]. Available: https://www.crowdstrike.com/cybersecurity-101/lateral-movement/

[38] 'Phishing accounts', BitDefender. Accessed: Nov. 27, 2023. [Online]. Available: https://www.bitdefender.com/blog/hotforsecurity/100-000-hacked-chatgpt-accounts-up-for-sale-on-the-dark-web/

[39] Crowdstrike, ‘Cyber big game hunting’, Cybersecurity 101. Accessed: Sep. 18, 2023. [Online]. Available: https://www.crowdstrike.com/cybersecurity-101/cyber-big-game-hunting/

[40] *Maersk ransomware*. Accessed: Nov. 27, 2023. [Online]. Available: https://www.digitalguardian.com/blog/cost-malware-infection-maersk-300-million

[41] Gartner, ‘PAM Tools’. Accessed: Sep. 25, 2023. [Online]. Available: https://www.gartner.com/reviews/market/privileged-access-management

[42] IBM, ‘What is IT infrastructure?’ Accessed: Oct. 31, 2023. [Online]. Available: https://www.ibm.com/topics/infrastructure

[43] X. Huang, Z. Guo, and M. Song, ‘FGLB: A fine-grained hardware intra-server load balancer based on 100 G FPGA SmartNIC’, *Int. J. Netw. Manag.*, vol. 32, no. 6, p. n/a, 2022, doi: 10.1002/nem.2211.

[44] Klaus. Schmidt, *High availability and disaster recovery : concepts, design, implementation*. in Advances in information security ; 22. Berlin ; Springer, 2006.

[45] A. Limited, *ITIL® 4*. London: The Stationery Office Ltd, 2020.

[46] C. Agutter, ‘10.3 Incident Management’, in *ITIL® Foundation Essentials - ITIL 4 Edition - The Ultimate Revision Guide*, IT Governance Publishing. [Online]. Available: https://app.knovel.com/hotlink/pdf/id:kt011XZTB1/itil-foundation-essentials/incident-management

[47] National cyber security centre, *Incident management*. Accessed: Nov. 30, 2023. [Online]. Available: https://www.ncsc.gov.uk/collection/10-steps/incident-management

[48] ‘ISACA Glossary’. Accessed: Oct. 15, 2023. [Online]. Available: https://www.isaca.org/resources/glossary

[49] ENISA, ‘IT continuity’. Accessed: Dec. 12, 2023. [Online]. Available: https://www.enisa.europa.eu/topics/risk-management/current-risk/bcm-resilience

[50] IBM, ‘Disaster Recovery’. Accessed: Oct. 16, 2023. [Online]. Available: https://www.ibm.com/topics/disaster-recovery

[51] M. J. Haber, ‘Break Glass’, in *Privileged Attack Vectors: Building Effective Cyber-Defense Strategies to Protect Organizations*, M. J. Haber, Ed., Berkeley, CA: Apress, 2020, pp. 189–202. doi: 10.1007/978-1-4842-5914-6_13.

[52] ‘Load balancing with HAProxy’. Accessed: Nov. 12, 2023. [Online]. Available: https://severalnines.com/resources/whitepapers/mysql-load-balancing-with-haproxy/

[53] ‘Load balancing’. Accessed: Dec. 11, 2023. [Online]. Available: https://www.nginx.com/resources/glossary/load-balancing/

[54] W. He, Z. J. Zhang, and W. Li, ‘Information technology solutions, challenges, and suggestions for tackling the COVID-19 pandemic.’, *Int. J. Inf. Manag.*, vol. 57, p. 102287, Apr. 2021, doi: 10.1016/j.ijinfomgt.2020.102287.

[55] Paul Cichonski, Tom Millar, TimGrance, and Karen Scarfone, *Computer Security Incident Handling Guide*, vol. Revision 2. in NIST Special Publication 800-61, vol. Revision 2. NIST. Accessed: Aug. 12, 2023. [Online]. Available: http://dx.doi.org/10.6028/NIST.SP.800-61r2