jamk

Privileged access management model for a managed service provider

Heikki Tuononen

Master's thesis
May 2023
Master's Degree Programme in Information Technology, Cyber Security

jamk | Jyväskylän ammattikorkeakoulu
University of Applied Sciences

jamk

Kuvailulehti

Tuononen, Heikki

Privileged access management model for a managed service provider

Jyväskylä: Jyväskylän ammattikorkeakoulu, toukokuu 2023, 56 + 4 sivua

Master’s Degree Programme in Information Technology, Cyber Security. Opinnäytetyö, ylempi AMK.

Julkaisun kieli: englanti

Julkaisulupa avoimessa verkossa: kyllä

Tiivistelmä

IT-palveluntarjoajat operoivat ylläpitäjinä useissa asiakasympäristöissä ja tarvitsevat ylläpitäjiensä käyttäjien- ja pääsynhallintaan hyvinmääriteltyjä ja turvallisia menetelmiä. Palveluntarjoajien erikoisuus pääsynhallinnan kannalta on se, että ylläpitoyteydet kohdistuvat usein palveluiden hallintaliittymiin ja niissä käytetään ylläpito-oikeuksin varustettuja tunnuksia. Ylläpitoyhteyksiin kohdistuva tietomurto voi altistaa asiakkaiden koko infrastruktuurin väärinkäytöksille ja palveluntarjoajan riskinä on paitsi vahingonkorvausvastuu, myös vakava mainehaitta.

Ensisijaisena tehtävänä oli työn tilaajan, Telia Cygate Oy:n, olemassaolevaa pääsynhallintamallia pohjana käyttäen luoda uusi, parannettu malli ja keskittyä erityisesti pääkäyttäjien oikeuksien- ja pääsynhallintaan. Toissijaisena tavoitteena oli tuottaa uudistetusta mallista kattava dokumentaatio, joka sisältää suunnitteluratkaisujen perustelut ja toimii perehdytyksenä aihepiiriin. Lisäksi dokumentaation suunniteltiin toimivan lähtökohtana pääsynhallintamallin jatkokehitykselle.

Dokumentaatiota ja erityisesti sen johdanto-osuutta varten tehtiin kirjallisuusanalyysi. Organisaation tietoturvavaatimukset pääsynhallintajärjestelmälle kartoitettiin yrityksen dokumentaatiota käyttäen. Vaatimuskartoituksen pohjalta luotiin mukautettu tarkastuslista järjestelmien vaatimustenmukaisuuden arviointiin. Sekä olemassaoleva että ehdotettu järjestelmä arvioitiin tarkastuslistaa käyttäen.

Tarkastuslistan perusteella saadut tulokset osoittivat, että ehdotettu malli parantaa pääsynhallintajärjestelmän vaatimustenmukaisuutta. Kirjallisuusanalyysin pohjalta luotu johdanto-osuus avasi oikeuksien- ja pääsynhallinnan keskeiset käsitteet ja kuvasi pääkäyttäjien oikeuksien- ja pääsynhallintajärjestelmien erityispiirteet. Ehdotetun pääsynhallintamallin dokumentaatio tarjosi pohjan mallin toteutukselle ja jatkokehitykselle, jota käsiteltiin myös johtopäätöksissä.

Avainsanat (asiasanat)

pääkäyttäjäoikeuksien hallinta, identiteetinhallinta, palveluntarjoajat

Muut tiedot (salassa pidettävät liitteet)

jamk

Description

Tuononen, Heikki

Privileged access management model for a managed service provider

Jyväskylä: JAMK University of Applied Sciences, May 2023, 56 + 4 pages

Master's Degree Programme in Information Technology, Cyber Security. Master´s thesis.

Permission for open access publication: Yes

Language of publication: English

Abstract

Managed service providers operating as administrators in multiple customer environments need well-defined and secure means of providing access to their employees. The distinctive characteristic of a service provider's access to customer resources is that the access is provided to the management plane using highly privileged credentials. A breach in such a scenario severely compromises the infrastructure of the customer organization. The service provider, in turn, is liable to compensation for damages and loss of credibility.

The main task was to use Telia Cygate's existing access management model as a basis to develop an improved model and focus especially on privileged access management. An additional objective was to provide thorough documentation of the novel model including the reasoning behind the design decisions to improve the administration of the solution. Additionally, the documentation was aimed to offer an introduction to identity and access management in general, and to provide a starting point for extended development of the privileged access management solution.

To introduce the concepts related to privileged access management, a literary review was conducted. Organizational requirements for information security of the solution were gathered using the company documentation. Based on the requirements, a custom checklist of requirements was developed to function as a meter of compliance for the assessed solutions. The existing and suggested access management solutions were described and assessed using the custom checklist.

Assessments using the custom checklist indicated that the suggested model would improve compliance with organizational requirements. The documentation resulting from the literature review introduced the core concepts of identity and access management and described how privileged access management systems function. The description of the novel management solution served as basis for implementation and for the continued development, which was further discussed in the conclusions.

Keywords/tags (subjects)

privileged access management, identity and access management, managed service providers

Miscellaneous (Confidential information)

Contents

Acronyms **3**
1 Introduction **5**
2 Methodology **6**
2.1 Study goals 6
2.2 Research methods 7
2.3 Data collection and analysis 9
2.4 Reliability 11
2.5 Ethical considerations 12
3 Identity and access management **13**
3.1 Digital identity and accounts 13
3.2 Authentication and authenticators 16
3.3 Authorization and access control 20
3.4 Digital identity threats and mitigations 22
3.5 Identity and access management and governance 28
3.6 Privileged access management systems 29
4 Description and analysis of the current PAM model **32**
4.1 Network environment 32
4.2 Organizational requirements 35
4.2.1 Documentation review 35
4.2.2 Use case survey 36
4.2.3 Check-list for requirements 39
4.3 The current PAM model 40
4.3.1 Description 40
4.3.2 Compliance review 43
5 The proposed PAM model **44**
5.1 Description 44
5.2 Compliance review 47
5.3 Implementation 49
6 Conclusions **50**
6.1 Results and reliability 50
6.2 Discussion and further development 51

References **53**

Appendices **57**

Appendix 1. Organizational requirements checklist 57

Appendix 2. Current PAM solution compliance 59

Appendix 3. Proposed PAM solution compliance 60

Figures

Figure 1 Main steps of qualitative research 8

Figure 2 Concept map of use case 2.0 elements 10

Figure 3 Account creation process 14

Figure 4 Authentication process 16

Figure 5 Federation relationships 20

Figure 6 Role hierarchy for server management 23

Figure 7 Identity governance model 28

Figure 8 Components of a privileged access management system 30

Figure 9 A simplified MSP network diagram 33

Figure 10 Simplified system function of a PAM solution 37

Figure 11 Current management model 41

Figure 12 The proposed management model 46

Tables

Table 1 List of all detected use case parameters 38

Acronyms

AAA	Authentication, authorization, and accounting
ACA	Advanced control and audit
AD	Active Directory
AAL	Authenticator assurance level
ACL	Access control list
CSP	Credential service provider
DAC	Discretionary access control
IBAC	Identity-based access control
IAM	Identity and access management
IG	Identity governance
IT	Information technology
IdP	Identity provider
IETF	Internet Engineering Task Force
IoC	Indicator of compromise
ISMS	Information security management system
ISO	International Organization for Standardization
ITSM	Information technology service management
ITU	International Telecommunications Union
MAC	Mandatory access control
MFA	Multifactor authentication
MITM	Man-in-the-middle
MSP	Managed service provider
NIST	National Institute of Standards and Technology
LDAP	Lightweight directory access protocol
PAM	Privileged access management
PAW	Privileged access workstation
PIM	Privileged identity management
PIN	Personal identification number
PKI	Public key infrastructure
RBAC	Role-based access control
RFC	Request for comments

RP	Relying party
SAML	Security assertion markup language
SIEM	Security information and event management
SoA	Statement of applicability
SOC	Security operations center
SP	Service provider
SSO	Single sign-on
TCSEC	Trusted computer system evaluation criteria
TFTP	Trivial file transfer protocol
TLS	Transport layer security
TOTP	Time-based one-time password
UBA	User behavior analysis
VDI	Virtual desktop infrastructure
VPN	Virtual private network

1 Introduction

Information technology (IT) services are a highly specialized branch of industry requiring substantial investments in infrastructure and skilled labor. Building datacenters with computing and networking capacity is expensive as is keeping technical specialists on the payroll. For companies that regard IT services as a core competency that they can do better than the competitors, the expenditure is a natural operating cost. Organizations not specializing in IT services might also elect to run their internal IT departments for compliance, confidentiality, or strategic reasons.

Organizations that do not consider producing IT services as their core activity, can view IT expenditure as an overhead that should be minimized, and the value of the service as hard to quantify. The traditional view has been that a specialized service vendor can leverage the economies of scale and technical expertise to provide better quality of service at lower cost. Over time motivation for outsourcing has developed from straightforward cost management to forming strategic alliances with service providers to enable more agile business development (Dibbern et al., 2004, pp. 7-8). Outsourcing business has become a major industry, and according to Grand View Research, the global IT service outsourcing market was estimated to be close to 500 billion Euros in 2019 and is expected to grow 7,7 % annually until 2027 (Grand View Research, 2022).

Managed service providers (MSP) manage their security posture like any other IT organization with a few distinctive elements. One special feature in the managed service provider business is that the security posture of the MSP directly affects the security of its customers. Customers trust their critical infrastructure in the care of the service provider and expect them to adhere to the industry best practices. Another special feature is that the specialists managing customer systems need to use administrative accounts that have greater control over the assets than normal user accounts. The administrative access also often targets the management plane of systems instead of the data plane used by normal users. This combination increases the severity of any breach. The greater privilege should mean that the accounts and their use are controlled and monitored more strictly. Getting access to one privileged account is a very attractive target to a malicious user, and an MSP environment that uses numerous administrative accounts to manage multiple customers is an even more lucrative target for an attack.

The importance of managing user identities and access is well-known since the early days of computing, multi-user time-share systems, and the first report on computer security (Ware, 1970). Several software and hardware vendors have products available to manage this specific area with different feature sets. For instance, Gartner, a well-known IT analyst and consulting company, lists 71 identity governance and administration products (Gartner, 2023). Since the privileged accounts pose the greatest risk to an organization, specialized tools are available for securing and managing privileged identities. Implementation scenarios in different organizations vary significantly so there is no silver bullet available, but the topic offers great opportunities for research and optimization. Again, MSP environments provide an interesting angle into the topic as service providers must operate in multiple identity domains and still be able to keep an indisputable audit trail of its users' actions.

The target audience for this work is the company commissioning the work and its architectural team. The aim, however, is to make the report general enough to be applicable to any complex ICT environment that needs to provide access to protected resources. The following chapters reviewing identity and access management concepts should also be usable for anyone interested in identity and access management (IAM) and privileged access management (PAM) in general, and the list of references should provide a good base for further information gathering in the subject.

2 Methodology

2.1 Study goals

The commissioner of the work, Telia Cygate, is a managed service provider i.e., a company that provides its customers with information technology services such as data center hosting, server management, and cyber security consulting (Telia Cygate, 2023). This thesis work is a part of a project that will implement an improved access management system and guidance. The outcome of this thesis should be a more secure access management model for accessing internal production systems as well as customer servers and devices. The model will cover all the use case scenarios that the organization has for its internal users and partners. Emphasis of the proposed model is in technical configuration of systems, but related policies and processes are also evaluated where applicable. A special focus is on implementing a PAM solution in a pre-existing environment. Additionally, the resulting documentation should serve as a useful introduction into privileged identity

and access management and provide tangible benefits of PAM systems to motivate the investment in software and labor.

Stakeholders in this work are business, employees, and the customers of the commissioner company. The main perspectives for evaluating the results of this work are the viewpoint of a system architect and the viewpoint of a system administrator. The architect view focuses in developing a scalable access model that can be applied directly into the current production environment, and which provides improved manageability and security. The system administrator view, on the other hand, concentrates on the usability of the proposed solution. The proposed solution should not add complexity to the work of a system administrator.

Currently, the access management solution does function but lacks a comprehensive documentation and is not fully centralized. An important part in the proposed version is its documentation of technical decisions and their rationale. Business requirements, planning decisions, perceived use cases, and selected controls are recorded diligently to provide a good understanding of the system. The documentation also establishes a base for future development.

The primary research question is whether the proposed model fulfills the organizational requirements. In this work, the organizational requirements are collected from multiple sources to form a checklist that can be then used to verify the validity of the proposed solution. The same checklist will also be used to compare the compliance rating of the existing and proposed solutions. This comparison produces an answer to the secondary research question: how does the novel approach improve security?

2.2 Research methods

The thesis work is a part of a larger internal development project in the commissioning company. The project requirements are directly sourced from the business needs of the company and the expectation is that the resulting work is based on verified knowledge in information security. Identifying the relevant trustworthy knowledge on the subject is an important goal in itself and makes the results interesting in a more general level. This description of the planned work fits well the definition of research-based development (Toikko & Rantanen, 2009, p. 22).

The main research and development subject, privileged access management system, is a complex entity consisting of several technical subsystems and different policies and processes, and the related research data is mostly textual in the form of documents. The research and data interpretation methods applied are therefore mostly qualitative instead of quantitative (Saunders et al., 2019, p. 175). The outline of qualitative research, as illustrated in Figure 1, also serves the purpose of this study well. Especially important parts of the process are the feedback loops linking back to step 4. As some of the data is likely unknown prior to the literary review and the information gathering of organizational needs, collection of further data will be inevitable during the project.

The expected result of the literary review is a collection of sources that provide a basis for theoretical knowledge in identity and access management and privileged access management. The secondary objective of the review is to provide a reference library for future development effort. Collection of organizational needs is expected to produce a concrete list of requirements that the outcome of the development must fulfill.

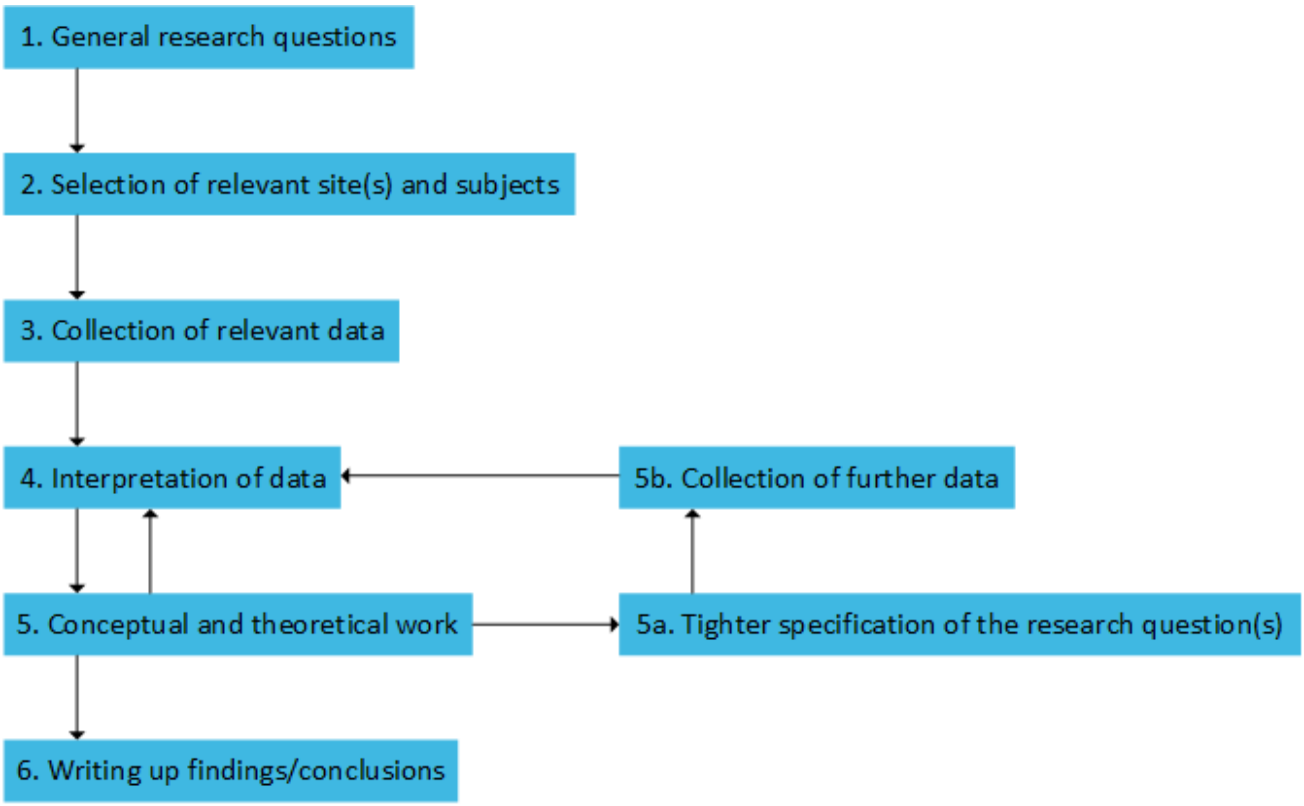

Figure 1 Main steps of qualitative research (adapted from Bryman, 2012, p. 384)

2.3 Data collection and analysis

The first task in the project is to find out what sort of theoretical knowledge, research, and industry standards exist in the domain. It is already known that the area is covered with numerous vendors offering products for privileged access management. Possible information sources are searches in article databases, library inventories, and the reference lists of other thesis works. Another important information gathering task is to enumerate the organizational requirements for the solution. This is done using publicly and internally available information security policies of the company and other relevant security documentation. Interview of employees working in architectural or security roles are also possible.

The theoretical part of this work uses the publications of National Institute of Standards and Technology (NIST) as its reference in concepts and glossary whenever possible. The main reasoning in this selection is to provide a coherent presentation of the subject matter by using one authoritative source. NIST also has a comprehensive list of publications freely available, and as a part of the U.S. Department of Commerce, it should be a reliable source of information. Complementary sources are used to verify the NIST definitions and to offer more holistic view of the subject.

Defining the functional requirements for a privileged access management system is identified as the third distinct data collection task in this project. The organizational requirements should outline the high-level design of security controls whereas the functional requirements should produce a more detailed list of technical aspects the design must deliver. The functional requirements for privileged access management are collected using use case scenarios. A singular use case is a description of how a certain user uses a system to achieve a certain goal focusing on how the system functions. A complete set of use cases lists all the different ways a system is used, and thus shows the value provided with the system. From the point of view of system design, use cases are a method of system development that helps the architect to understand the ways of using the system and to design a system that supports the users. A documented list of use cases describes clearly what a system does, and by intentional omission, it also describes what the system does not do (Jacobson, Spence & Bittner, 2011, p. 4).

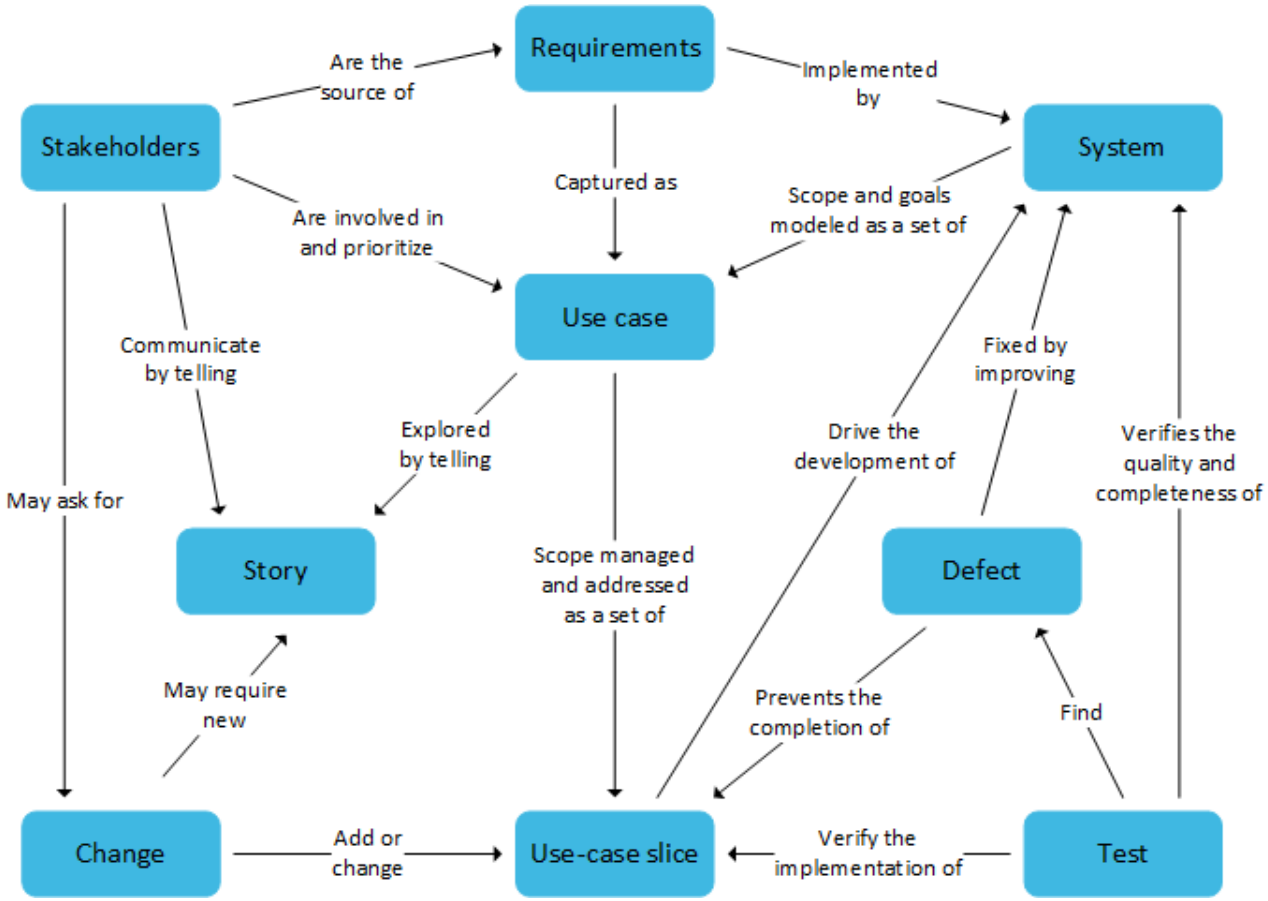

Figure 2 Concept map of use case 2.0 elements (adapted from Jacobson, Spence & Bittner, 2011, p. 13)

Figure 2 shows elements of use case 2.0 practice and their relations as described by Jacobson, Spence & Bittner. Story and case slice, along with use case, are the key elements for capturing system requirements. Story is a stakeholder's detailed description of how to achieve a particular goal and how to handle possible issues focusing on the result and benefits of the result. Use case slice is a concept that helps to break down a use case to smaller pieces to simplify development and testing effort.

For this work, a top-down approach is used. For the top-down approach, the workflow is the following:

1. Identifying the use case
2. Outlining the steps in the basic flow
3. Identifying possible alternative flows based on the basic flow.

As opposed to designing an entirely new solution, this project benefits from having an existing solution in production use. The already deployed system gives a good baseline of expected product

functionality and reveals opportunities for improved functionality. Having an existing production system also provides a great starting point for identifying different use cases including the alternative flows. As some system components will change, it is reasonable to expect that some features will change as well, affecting the use cases.

2.4 Reliability

Compared to quantitative research, assessing the reliability and validity of qualitative research is a complex and somewhat arbitrary task as the research process relies more on interpretative assumptions. In general, reliability means that the research is consistent and can be repeated. Valid research is generalizable, provides accurate analysis and results, and uses appropriate measures (Saunders et al., p. 213 - 214).

In this project there are two distinct points where evaluation of reliability is especially needed. The first point is after the information collection and analysis phase. At this point the deliverables are an introduction into access management sourced from the literary review, the list of requirements for an access management solution sourced from organizational requirements, and the functional requirements sourced from a use case analysis. This work strives to ensure the reliability of information gathering by selecting good quality sources that constitute a representable sample of available knowledge in the area.

Since literary sources in the field are numerous, identifying the essential sources is vital. Effective use of search engines and thorough investigation of reference lists of books and peer-reviewed articles are expected to narrow down the key sources. As this work mostly deals with basic concepts of access management, selecting a representative set of sources should not be an issue since the concepts are generally well normalized and contradictory or conflicting information or definitions are not expected within the literature. It is reasonable to expect that even by selecting an entirely different set of sources, a repeated study could arrive at similar results with few differences. The most notable source of unreliability is possible omission of an informative or a contradictory source or sources.

In the case of gathering the organizational information, the data is much scarcer and more concentrated. As the strategy is to use only public sources wherever possible, the data should be

easily collected and analyzed. Here the sample size is approximately the size of the entire data, trustworthiness of the source is high, and thus reliability reduces to the conduction of the analysis. Given the relatively small amount of data for analysis, the reliability is expected to be high.

The use case analysis uses actual user actions including the author's experience of using and administering the existing solution as input data. Together with the process outlined in chapter 2.3, this should lead to a comprehensive description of potential use cases at the time of the writing. New features in the selected software can potentially enable new use cases, especially over time as the software is developed further, but that is not in the scope of this work.

The second point, where a reliability assessment is in order, is when the novel access management solution is designed. The reliability of the proposed model naturally relies on the reliability of the list of organizational requirements, which is high, as described above. Repeatability is a difficult factor to evaluate here, since the design process is not entirely deterministic, but consistency should be ensured by documenting and justifying the design choices.

2.5 Ethical considerations

This work is focuses on technical systems and their development that are researched using literary sources and firsthand knowledge of the subject environment. During the information gathering the study subject should not pose many chances for violating the tenets of ethical behavior in scientific research as described by Bhattacherjee (2012, p.137). In this sort of study, the analysis and reporting sections are the most potential ones for compromise. According to Bhattacherjee (2012, p.139), openness and honesty are the driving force of science, and the best way for researchers to help the progress is to fully disclose even the unexpected or negative findings. This study complies with that description and reports the results truthfully.

Since the work is commissioned by a company planning to use the results in its production environment, the level of detail in the report must be carefully balanced. Some data, such as collection of organizational requirements for PAM, demand the use of company's internal information. Presenting the findings in too general level makes the result less interesting for the reader and less usable for the commissioner. Alternatively, too detailed reporting might create security risks for the commissioner.

3 Identity and access management

The basic concepts referenced throughout this work are user identities, authentication, and authorization. Building on the basic concepts, an introduction to access management and access controls to networks and systems is presented. These elements are important building blocks of information security and there are numerous frameworks, methods, and products available for managing them. However, there are few publicly available models encompassing all these elements and providing enough detail to implement a solution in a pre-existing production environment. The lack of the readily available solution is the key motivation that set this development in motion. To provide some background into the subject and to put the observed solutions into context, this chapter describes the basic concepts related to privileged access management.

3.1 Digital identity and accounts

Information security is about protecting information against access without permission, alteration, or destruction. Digital information systems are predominantly multi-user systems that have numerous features and store vast amounts of information. To protect the information, users of a system must be identified to assign correct level of access to features and data. Reliable identification also protects users from identity theft (Windley, 2023, Chapter 11).

In information security, subjects are represented by their digital identity when accessing protected resources. Subject does not need to be a real-life subject such as a person or user, but can refer to a system, a service, or a function that has no physical characteristics. A subject can be also referred to as an entity, an applicant, a claimant, or a subscriber depending on the context. According to NIST SP 800-63-3, a digital identity can be defined as “an attribute or set of attributes that uniquely describe a subject within a given context”. The context is important as the same identity may not be unique or valid in different contexts and the same subject can have different identity in different contexts (Grassi, Garcia et al., 2017, p. 47).

The recommendation ITU-T Y.2720 for next generation networks' identity management framework categorizes the information content of a digital identity to contain the following groups (International Telecommunications Union, 2009, p. 5):

- Identifiers are unique data or labels used to indicate subjects and their attributes within a context. For example, usernames, email addresses, and employee numbers (National Institute of Standards and Technology, 2020, p. 403).
- Credentials are objects or structures that connect identities and possibly attributes, represented by identifiers, to authenticators that are owned and controlled by subjects. For example, a username and an associated password (National Institute of Standards and Technology, 2020, p. 400).
- Attributes are data describing subject's characteristics or quality. For example, full name, roles, privileges, and location (Grassi, Garcia et al., 2017, p. 40).

In an organizational context, a natural person should have only one digital identity that is tied to a unique identifier such as an employee number or a social security number. Typically, digital identities are represented as accounts. One identity can have multiple accounts that represent the person's digital identity and different roles. For a given system an account and its relation to an identity is usually straightforward as the account represents one identity and one user. Conversely, a single user can have tens of accounts for different systems and even multiple accounts for a single system to represent different roles and permissions (Haber, 2020, pp. 10-11).

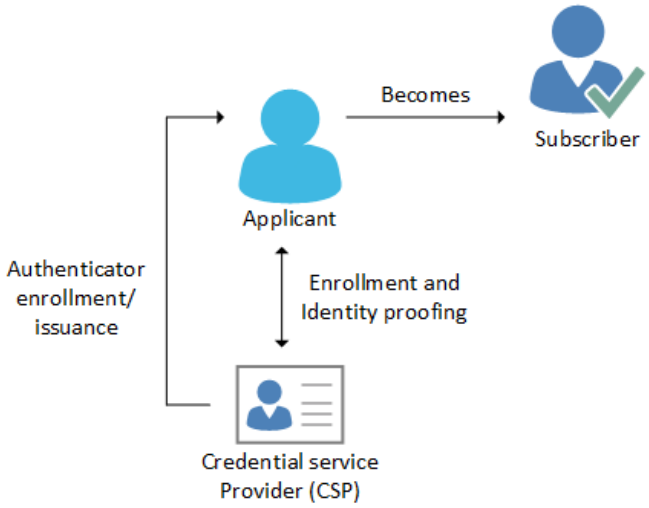

Figure 3 Account creation process (adapted from Grassi, Garcia et al., 2017, p. 10)

Figure 3 shows the account creation process. The process starts with enrollment where a requester or an applicant contacts a credential service provider (CSP). Depending on the application,

the applicant is assigned an identifier or allowed to choose one, such as an email address. In addition to the chosen identifier, the applicant may need to provide extra attributes (Windley, 2023, Chapter 11).

Identity proofing is a part of the account provisioning process that ensures that the account is linked to a specific requester. Depending on the account type and usage, the reliability of the proofing method varies notably. A free web service such as a news site registration typically only requires a verification that the requester has access to the email that was provided as an identifier during the enrollment. By contrast, opening a bank account requires stronger verification such as presenting a government issued proof of identity (Windley, 2023, Chapter 11). Successful identity proofing makes the applicant a subscriber and credential and authenticator(s) are established. The CSP maintains the credential and the subscriber maintains his or her authenticator(s) (Grassi, Garcia et al., 2017).

This study focuses specifically on user accounts, since they are most interesting regarding privileged access management. User accounts can be divided into two main categories: local accounts and centralized accounts. Local accounts are, as the name suggests, local to the resource where they can be used for logging in. Local accounts are supported by most systems and their provisioning requires no extra infrastructure. On the downside, local account management is tedious unless there is an identity management solution that can be used for provisioning, secrets management and other account administration tasks. Local accounts are also often duplicated across systems on purpose or inadvertently, which increases the potential attack surface if an account is compromised (Haber & Rolls, 2020, p. 32).

Centralized accounts leverage a directory service to store the account details. For enterprise accounts, Microsoft Active Directory (AD) is perhaps the most common solution (Microsoft, 2022a). Directory-based accounts are more user-friendly since one account can be used for multiple resources. Centralized management also decreases the administrative overhead as user management can be done using a single administrative interface (Haber & Rolls, 2020, pp. 32-33). Accounts can have group membership as their parameter, which is an especially useful feature of centralized accounts. Groups make management of large number of accounts easier as permissions can be granted to groups instead of individuals (Chapple, 2021, Chapter 1).

3.2 Authentication and authenticators

Authentication process verifies the identity of a subject. Usually this is done to provide access to a protected resource (Grassi, Garcia et al., 2017, p. 41). Verifying an identity also establishes confidence that the subject has the attributes bound to the identity or the account that represents the identity (Grassi, Fenton et al., 2017, p. 2). It must be noted that authentication alone does not determine the subject's access to resources, although systems often obfuscate this distinction from the user (Haber & Rolls, 2020, p. 12-13). To put it another way, authentication is a mechanism that blocks outsiders from accessing a protected resource (Windley, 2023, Chapter 11).

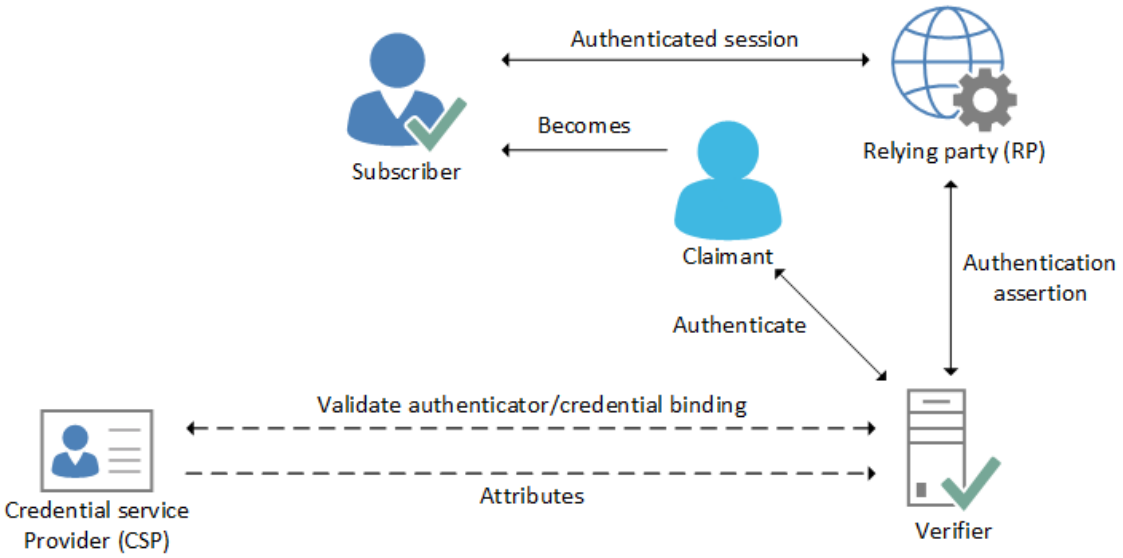

Figure 4 Authentication process (adapted from Grassi, Garcia et al., 2017, p. 10)

Figure 4 shows the digital authentication process. In the authentication process a claimant asserts an identity and demonstrates control of an authenticator to the verifier using an authentication protocol. The verifier validates the credential by itself or with the CSP (Grassi, Garcia et al., 2017, p. 14). A successful authentication process asserts the claimed identifier and possibly other identity information to the relying party (RP), which usually has the protected resources that the subscriber requests access to (Grassi, Fenton et al., 2017, p. 2). The authenticated identity, or the subscriber, can then be linked to a user account in the system that is the actual vehicle used for gaining access to secured resources (Haber & Rolls, 2020, p. 22). An entity that can be authenticated is also known as a security principle (Microsoft, 2022b).

Authenticity of a subject is determined using one or more factors or authenticators. The strength and robustness of authentication can be increased by adding factors. Three classic categories of factors a subject can provide are (Grassi, Garcia et al., 2017, p. 12):

- Something you know. A secret such as a password that is known to the subject only.
- Something you have. Usually, a physical object such as an ID card or an application registered to a certain phone number.
- Something you are. A characteristic of a subject such as facial features or a fingerprint.

The something you know factor is ubiquitous as passwords and personal identification numbers (PIN codes) are widely used in securing access to all sorts of accounts (Windley, 2023, Chapter 11). Passwords are the most common example of shared secrets or symmetric keys. A password is a memorized secret whereas symmetric cryptographic keys are usually stored in software or hardware the subject controls. Although a long and complex password is inherently like a cryptographic key, system generated keys tend to be safer. This is due to the human tendency to pick less complex passwords for memorization. As opposed to shared secrets, public key pairs do not require transfer of a secret between the subject and the verifier. Instead, the verifier acquires the public key associated with the subject's identity and uses an authentication protocol to verify that the public key matches the private key presented by the subject (Grassi, Garcia et al., 2017, pp. 12-13).

The something you have factors prove authenticity by verifying that a subject has possession of a token or a resource that was bound to the identity earlier, usually during identity proofing. Phone numbers, mobile apps, email addresses, application recovery codes, and hardware tokens are examples of digital methods. The mechanism of verifying the possession depends on the token type. Phone numbers and emails can be verified via codes relayed by text messages or email messages and a mobile app or a hardware token can generate time-based one-time passwords (TOTP) (Windley, 2023, Chapter 11).

The something you are factor is an inherence factor describing a unique feature of a person. Also known as biometrics, this factor requires scanning a feature such as a fingerprint or facial features to create a data structure representing the feature. Subsequent scan results are compared to the data structure to determine authenticity. Biometrics is a trustworthy way of proving authenticity and more user-friendly than typing passwords or PIN codes, making it popular especially in mobile devices (Windley, 2023, Chapter 11).

Other attributes of the authentication event, such as source IP addresses and related geolocations, times of events, or source device identifiers can be used to evaluate the risk involved in the claimed identity, but they are not authentication factors (Grassi, Garcia et al., 2017, p. 12). Such attributes are examined further in chapter 3.4.

An authentication event that relies only on one factor is called single-factor authentication. In computer systems the single factor used to authenticate a given username or an identity is usually a password. The level of confidence of authentication can be increased by requiring more factors. Since a password can be stolen or intercepted multifactor authentication (MFA) is usually needed for strong authentication. MFA requires the use of at least two independent authentication factors from different factor categories. A typical example of combination of authentication factors in computer systems is a password that the user knows and a mobile phone application that the user has (Chapple, 2021, Chapter 1).

One framework for measuring the strength of an authentication mechanism is the Authenticator Assurance Level (AAL) as described in NIST SP 800-63-3 (Grassi, Garcia et al., 2017) and in NIST SP 800-63b (Grassi, Fenton et. al., 2017). The AAL model contains three levels AAL1 - AAL3, where a higher digit denotes higher confidence that the claimant controls an authenticator or authenticators that are bound to the target account. Accordingly, the three levels have different requirements for the number of authenticators, permitted authenticator types, reauthentication intervals, and other security controls.

The most important protocols providing authentication services in the context of this work are Lightweight directory access protocol (LDAP) and Kerberos. LDAP is a vendor-neutral protocol that is used for connecting to a X.500 directory service database and querying user information. LDAP specification is defined in IETF's RFC 4511 (Semersheim, 2006). LDAPS or LDAP over SSL/TLS is an important yet a non-standard addition to LDAP protocol to provide encrypted communication by establishing a TLS connection prior to exchanging any LDAP messages. LDAP is supported by many directories and applications, which enables interoperability between products and vendors (Windley, 2023). Active Directory is among the user directories that support LDAP and LDAPS (Microsoft, 2014).

Kerberos, as defined in IETF RFC 4120, is a protocol that uses a trusted third-party authentication service and symmetric-key cryptography to provide secure authentication for principals that are pre-joined members of the Kerberos realm (Neuman et al., 2005). Key distribution center (KDC) is the third-party acting as an authentication service and a ticket-granting service. The granted tickets are used for authenticating to services in the same realm. The principals do not need mutual trust between them but do trust the KDC. Active Directory domains are one common example of environments implementing Kerberos authentication (Chapple, 2021, Chapter 5).

Kerberos can be leveraged to implement single sign-on (SSO). SSO enables users to login once and then gain access to systems within the logon domain without typing their credentials again. Elimination of multiple accounts and passwords streamlines user experience, helps organizations to enforce the same password policies across all resources, and decreases the administrative overhead that comes with managing multiple accounts for each user. A possible disadvantage is that one compromised SSO account can be used to access multiple services (Chapple, 2021, Chapter 5).

Federation is an authentication concept that relies on a trusted third party across organizational boundaries. For instance, many web services use authentication services provided by global technology companies such as Facebook and Google. The party providing the authentication service is an Identity provider (IdP) and it creates and maintains trusted identity information of other entities. Service provider (SP) is the party trusting the assurance of the identity provided by the IdP (Chapple, 2021, Chapter 7). Service provider can also be referenced as the relying party (RP). The relationships between the subject, the IdP, and the SP are presented in Figure 5. The subscriber initiates the federated authentication process by requesting access to a service, which redirects the subscriber to authenticate using the IdP. After successful authentication the IdP redirects the subscriber back to the SP. The authentication status is communicated by the IdP to the SP either through redirects (front channel) or using a direct connection (back channel) (Grassi, Richer et al., 2017, pp. 8-9).

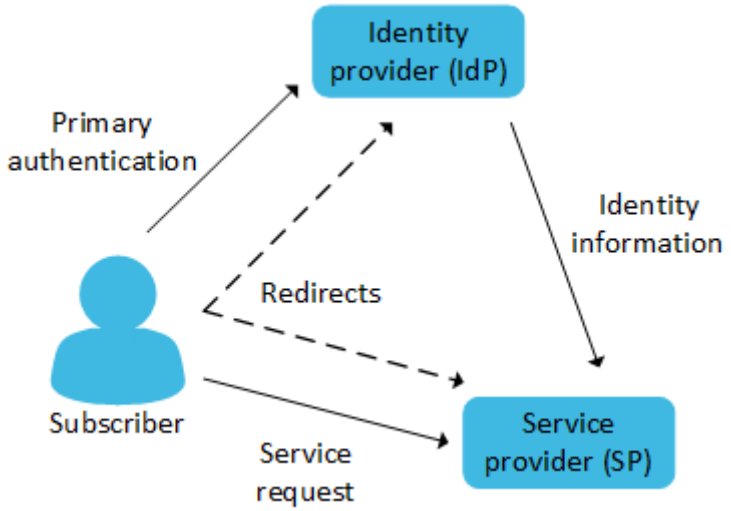

Figure 5 Federation relationships (Adapted from Grassi, Richer et al., 2017, p. 8)

The benefit of federation is that the users can use a single set of credentials to authenticate to multiple services. Security assertion markup language (SAML) and OpenID connect are common implementations of federated identity (Chapple, 2021, Chapter 7).

3.3 Authorization and access control

Authorization is the method of granting access to resources as specified by the access rights linked to an account. Authorization only takes place after successful authentication. The access rights of an account define what sort of functions the account can perform and can also explicitly deny some functions. The rights can be assigned in the application, operating system, in the supporting infrastructure, or in an external identity or a privilege management system (Haber & Rolls, 2020, pp. 13-14).

Access control formalizes the rules how a subject can interact with protected resources. More specifically, access control defines and grants rights, or privileges, pertinent to a target resource, to a subject that requests access (Chapple, 2021, Chapter 1). In other words, access control defines a set of authorizations that an account has. Thus, authorization and access control are often used interchangeably (Windley, 2023, Chapter 12). AAA is a well-known model or framework, that unites the three pillars of access control: authentication, authorization, and accounting. Accounting complements the two former pillars by keeping records of actions that authorized users perform (Chapple, 2021, Chapter 1).

The permissions that an authorization grants are naturally resource specific and depend on the interface that is available to a subject requesting access. For this study, the most relevant resource types are servers and services that are accessed using network connectivity.

Various methods can be applied to authorizing access to protected resources. Most access control decisions are based on the identity of the user and can thus be described as identity-based access control (IBAC). The discretionary access control (DAC) is a model dating back to 1983 and Trusted Computer System Evaluation Criteria (TCSEC) published in DoD 5200.28-STD, also known as the Orange Book (Department of defense, 1985). The model is still valid and implemented ubiquitously, for instance in filesystems such as Windows NTFS and Linux/Unix POSIX, using access control lists (ACL). In DAC model, access to objects is restricted based on the identity of subjects or their groups. Since the control is discretionary, a subject with certain permissions, such as ownership of the object, can delegate permissions to other subjects unless mandatory access control prevents it (Chapple, 2021, Chapter 5).

Mandatory access control (MAC) is also described in TCSEC. In MAC the objects are labeled according to the sensitivity of their information content and managed by administrators only. This contrasts with the DAC model where subjects can grant permissions. The application of MAC is mostly done in defense-related organizations where the subject's security clearance level must determine the access to resources (Chapple, 2021, Chapter 5).

Both DAC and MAC focus on the information content of the protected resource whereas role-based access control (RBAC) or nondiscretionary access control focuses on the user and user's role in the organization or within the system to grant access. An RBAC system considers the following three aspects (Chapple, 2021, Chapter 5):

- Role assignment: A subject must have a role assigned and active to take an action allowed for that role.
- Role authorization: The active role of a subject must be authorized for the subject ensuring that the subject cannot assign permissions for oneself.
- Transaction authorization: A transaction can be executed only if it is authorized for the active role of the subject.

RBAC is a concept of higher level of abstraction than DAC or MAC and requires more administrative work to implement. The advantage of RBAC is that, once implemented, the administration is simpler and requires less effort. For instance, an organization can have a base role for a certain

organizational unit such as IT. The base role can include the common permissions for that unit and additional permissions can be assigned using more granular roles such as a database server management role. (Chapple, 2021, Chapter 5).

RBAC helps to fulfill two important design principles in cyber security: the principle of least privilege and separation of duties, originally introduced by Saltzer and Schroeder (1975). The principle of least privilege states that a user should have only the permissions or access to perform the tasks required by her or his job. Separation of duties aims to protect systems from insider threat by dividing permissions in a manner that prevents a single user from compromising the system. An example of least privilege in an MSP environment would be a specialist who can login to servers of customer A to perform administrative tasks but cannot login to servers of customer B because managing customer B is not mandated by the specialist's job role. To extend the same example for the principle of least privilege, the access logs of servers should be shipped to a system where the specialist cannot alter or delete them to hide her or his actions (Chapple, 2021, Chapter 5). This is also a good example of how accountability augments gaps in access policy enforcement. Tamper-proof logging increases the risk of being held accountable for unauthorized activity and thus acts a deterrent (Windley, 2023, Chapter 12).

3.4 Digital identity threats and mitigations

A digital identity or a user account with credentials is an attractive attack target for any malicious actor. Often a compromised identity provides a shortcut into resources bypassing network-level defenses (Microsoft, 2022c). To rephrase and to stress the importance of the previous statement, one could say that "identity is the new perimeter", meaning that as services are moved from on premises to cloud, a traditional network perimeter is harder to establish and authenticating users becomes the principal protection method of people and systems (Windley, 2023, Chapter 11).

As a measure of importance of identity protection, Verizon's data breach investigations report for 2021 shows that around 40 % of data breaches included in the report involved the use of stolen credentials (Bassett et al., 2022, p. 15). Another recent study from f5 reports that although the amount of credential spill incidents is on the rise, the actual number of spilled credentials has decreased during the last five years (Vinberg & Overson, 2021, p. 8).

For this study, it is important to consider the implications of a compromise for different account types. Typically, the principle of least privilege is implemented assigning a group of accounts the relevant permissions that are required for a certain job role. This yields granular privilege levels from standard users to full administrators (Haber, 2020, p. 4). One example of such role hierarchy for server management is displayed in Figure 6. The standard user is a valid organization account but does not have the permission to login to the server, as this is not mandated by the job role. The server operator can login and use applications pertinent to the role but cannot make changes to the system configuration. The administrator has unrestricted access to the system and its settings, including controlling the permissions for other users. From the viewpoint of a malicious user, the account with the elevated permissions is usually the most interesting target.

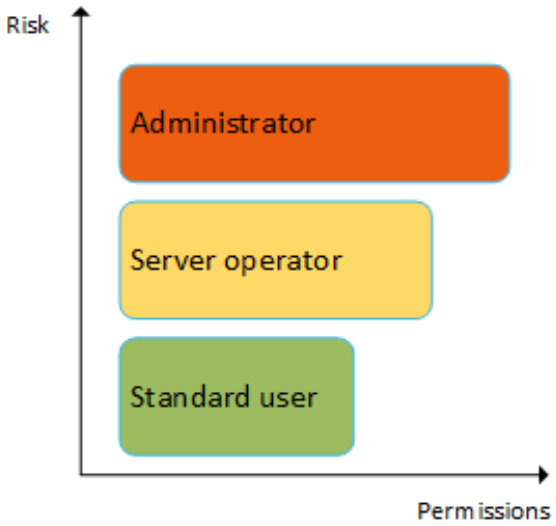

Figure 6 Role hierarchy for server management (Adapted from Haber, 2020, p. 9)

In some cases, an account with data plane privileges can be as valuable to an attacker as an account with management plane privileges. For example, a physician's account probably has access to patients' health information whereas the administrator managing the health care application might lack access to health data as it is not required for his or her job role (Haber & Rolls, 2020, p. 113).

Observing the risk levels in Figure 6 leads to one obvious conclusion: The more privilege an account has, the more protection the account requires. One good source for account and access protection methods and policies is NIST Special Publication 800-53 (National Institute of Standards

and Technology, 2020). Its access control chapter lists, among others, the following protection measures for general and privileged account management:

- Define and document the account types that are allowed and the account types that are specifically disallowed in each system.
- Define an approval process for provisioning accounts.
- Create an account life-cycle process for creating, modifying, and removing accounts and align it with personnel processes such as recruitment, transfers, and termination.
- Monitor the use of accounts and review account compliance.
- Administer privileged accounts using a role-based or an attribute-based scheme.
- Monitor the assignment of privileged roles or attributes and changes made to those roles or attributes.
- Enforce separation of duties where necessary and align account management accordingly.

Regarding the principle of least privilege, the same publication outlines the following guidance for access control:

- Authorize access to security functions and security-related information.
- Use non-privileged accounts for accessing non-security functions to limit exposure.
- Separate processing domains to allow more fine-grained access policies.
- Assign privilege only to defined roles or personnel.
- Review privileges to validate the need for them and reassign privileges to align with organizational policies and requirements.
- Monitor the use of privileged functions.

In the context of this work, one key point in these instructions needs further examination: The separation of use of non-privileged and privileged accounts. Administrative users need both account types, and it is imperative to compartmentalize the privileged account use cases. An example of a vendor approach to this issue is the privileged access strategy developed by Microsoft (Microsoft, 2023). The strategy advocates isolating the use of privileged accounts only to designated devices and intermediaries. This limits the use of privileged credentials to well-defined points, such as jump servers, minimizing the attack surface. Microsoft's privileged access strategy also outlines the use of privileged access workstations (PAW). PAW is a hardened desktop that offers functionality only for management tasks. Risky tasks such as web browsing or opening emails or PDF files are not possible, which thwarts many popular attack scenarios and techniques.

Another interesting classifying factor with account security is the division between local and directory-based accounts. Local accounts only provide access to one system and therefore a compromised account only compromises one system given that the account is unique. Local accounts with the same credentials in one environment, such as a common admin account for multiple servers, expose the environment to privileged attack vectors and lateral movement. Equally, a

compromised directory account can compromise multiple systems. A compromised privileged account that can modify accounts and directory settings jeopardizes the whole directory (Haber & Rolls, 2020, pp. 32-33).

A shared account is a special account type that needs to be considered separately. A shared account does not have a 1-to-1 relationship with a single user, but the account credentials can be accessible to multiple users. This leads to a lack of accountability as the actual user of an account cannot be determined. Although shared accounts are a known bad practice, they do exist and have legitimate use cases, for instance in providing multiple administrators access to a managed system. To mitigate the risks of shared accounts, an organization should monitor the access to shared credentials and the use of shared credentials. Shared credentials should also be rotated periodically to keep the accounts from expiring and to prevent credential leaks after employee changes (Haber, 2020, pp. 25-26).

As described in chapter 3.2, accounts can be protected with different authenticators categorized into what you know, what you have, or what you are. A password is something you know and despite its shortcomings, it is still the most used method of authentication to access an account. Password hashing is a method to prevent storing passwords in cleartext (OWASP, 2021). Unfortunately, even hashed passwords can be used in attacks using pass the hash technique (MITRE, 2021). Since cleartext passwords are in essence just short strings, they can be compromised in many ways (Chapple, 2021, Chapter 5):

- Social engineering techniques, such as phishing, deceive users to input their credentials to a system or form where the attacker can recover them in plain text.
- Brute-force attacks systematically try different passwords until a match is found.
- Dictionary attacks use words found in dictionary to try to guess passwords quicker than using a brute-force attack.
- Eavesdropping unencrypted network traffic, man-in-the-middle attacks (MITM) or using a system contaminated with trojans or keyloggers can expose credentials.

Reusing passwords makes the issues even greater, since data breaches can expose credentials and passwords that can be then used to access other systems in credential stuffing attacks. The conventional methods for increasing password security are enforcing a suitable length, complexity, and lifetime for passwords. Such measures can be tedious especially to users who do not rely on password managers to offload the burden of remembering complex passwords. However, these methods mitigate the effectiveness of brute-force attacks and decrease the time a stolen account

is usable. Then again, the validity of these methods is questioned in NIST Special Publication 800-63B, which in turn outlines the following best practices (Grassi, Fenton et. al., 2017, pp. 13-15):

- Password changes should not be required unless the password is suspected to be compromised. Repeated expiration promotes using weaker passwords and possibly writing them down.
- A minimum length of eight characters should be the only complexity requirement. Extra complexity makes remembering a password more difficult.
- Password filtering should be implemented to prevent the use of passwords found in dictionaries or password dumps.

Password filtering helps against users choosing weak passwords that are prone to dictionary attacks and brute-forcing. Eavesdropping and traditional MITM are often mitigated with transport encryption as this protection is built-in in most modern applications and protocols. Social engineering is probably the toughest risk to mitigate, since it requires training the users to detect fraud and creating verification methods to help users avoid fraudulent services.

Resetting a user's personal account password is a very common procedure. Depending on the account type the process of resetting varies. For organization accounts a call to the help desk is a typical method whereas online accounts often provide a self-service password reset. Before replacing the lost password, the user needs to be authenticated using another factor. The help desk can use a registered phone number and self-service can use security questions or access to a registered email. Regardless of the authentication method used, it is vital to note that the authentication process is only as secure as the process of account recovery (Windley, 2023, Chapter 11).

Some initiatives are in place to retire the use of passwords for authentication. For instance, Microsoft is already rolling out a passwordless way for customers to use their online services (Jakkal, 2021). In this case the user can opt to use a mobile app, Windows Hello, a security key, or a verification code sent to phone or email instead of a password. Microsoft is also part of the FIDO ("Fast IDentity Online") Alliance and recently published its commitment to support a joint effort with Apple and Google to create a common passwordless sign-in standard for consumers to use across devices and platforms (FIDO Alliance, 2022).

Physical tokens and devices represent something you have. Used alone, they can provide physical security like a token that is shown to a reader device and allows a worker to gain access to an office building. Combined with something you know such as a credential, a device or token provides

an extra factor of authentication. A malicious user that gains access to either the credentials or the token cannot use them alone for identity spoofing (Chapple, 2021). Authenticators proving possession are susceptible to attacks such as TOTP code phishing, theft of physical tokens, and even theft of authentication cookies using the infostealer malware (Smilyanets, 2022).

A recent targeted phishing attack against Cloudflare proved the efficacy of multifactor authentication, when implemented correctly. Although some credentials were stolen, no breach was detected since the company uses hardware keys. What is notable in this attack, is the real-time relay of credentials and TOTP codes to the attacker, which would have been a successful strategy against a company using TOTP codes in MFA implementation (Prince et al., 2022). Another example with contrary consequences is provided by Deloitte, which in 2017 experienced a major data breach after the loss of one piece of administrative credentials that was not protected with MFA. The admin account to the company's email system was leveraged for lateral movement giving access to several other systems resulting in at least six of Deloitte's customers being impacted (Hopkins, 2017).

The use of biometrics represents something you are. Common examples of biometrics are fingerprints and facial recognition. Like tokens and devices, biometrics can be used alone for single factor authentication or in combination with other mechanisms in MFA. Although physical features measured in biometrics are unique, spoofing is possible. Apple's implementation of a fingerprint reader in a mobile phone, Touch ID, was hacked less than 48 hours after its release using a thin latex film with replicated fingerprint grooves (Frank, 2013). Spoofing Apple's facial recognition system, Face ID, took a couple of weeks and was done using a mask made of stone powder enhanced with 2D images of the eye areas (Bkav, 2017).

In addition to the three authentication factor categories described above, additional factors might be available for evaluating the authenticity of authentication. A location factor can use account attributes or the detected usual geolocation of the user. A temporal factor can consider the static office hours defined in account attributes or detected usual activity times of an account. Combination of login locations and login times can detect possible fraudulent login attempts if a user tries to login from two geolocations within a time that does not make travel between the locations possible (Windley, 2023, Chapter 11).

One recent development involving identity and access control is zero trust. The key idea of zero trust is that there should not be an implicit trust based on the physical or network location or asset ownership. Zero trust represents a response to increased use of cloud services, bring your own device (BYOD), and remote users. Instead of protecting network segments with static firewalls, dynamic protection of users, assets, and resources is in focus (Rose et al., 2020, pp. 4-5). Zero trust is also an integral concept in Microsoft privileged access strategy, discussed earlier in this chapter (Microsoft, 2023).

3.5 Identity and access management and governance

As made apparent by the preceding chapters, digital identity and access are complex to manage yet crucial to the usability and security of digital systems. The previous chapter described some polices and technical measures for securing identities, and this and the following chapter will describe frameworks for systematic management of digital identities.

Identity and access governance is a broad discipline outlining the methods and processes for ensuring that the users have appropriate access to resources and that the organization can track who has access to what, how the access can be utilized, and whether the access is compliant with organization policies. The organizational view in identity governance (IG) is displayed in Figure 7 (Haber & Rolls, 2020, pp. 45-47).

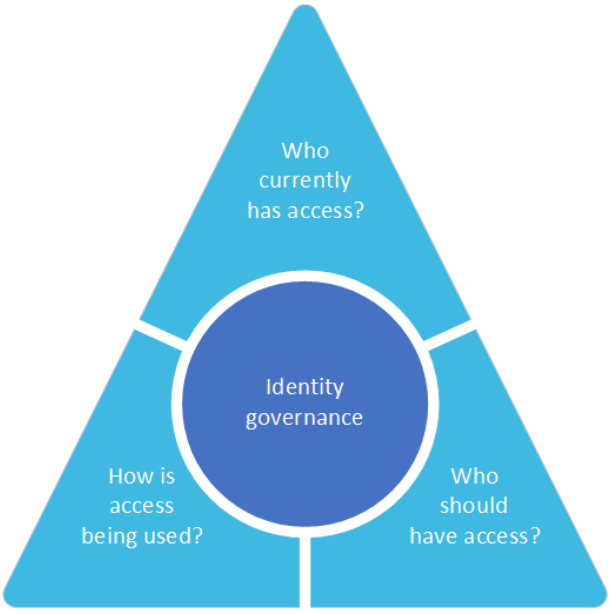

Figure 7 Identity governance model (adapted from Haber & Rolls, 2020, p. 46)

Who currently has access describes the existing situation, which is the first step of implementing an identity management system and important status information once a management system is operational. Who should have access describes the desired state, which is typically sourced from organizational policies. The final piece, how access is being used, provides important feedback on the process quality. Misused, unused and incorrect access should be tracked to improve access management (Haber & Rolls, 2020, pp. 46-47).

Identity and access management (IAM) is a technical solution that enables identity governance. In practice identity management and access control are often interwoven and implemented together. IAM solutions provide authentication services, access management, and account life-cycle management including provisioning, administering, and revoking user accounts (Chapple, 2021, Chapter 1).

Even a small organization benefits from identity governance. A modest number of resources on premises and in cloud with different privilege levels of accounts and a handful of employees, contractors, and partners are enough to create a web of privileges that is difficult to manage manually. Asset inventory, status visibility, and account automation significantly improve the organization's control over its resources (Haber & Rolls, 2020, pp. 47-48).

3.6 Privileged access management systems

The acronym PAM is used as a reference to privileged access management and privileged account management. Privileged identity management (PIM) and privileged user management (PUM) are also terms that are often used interchangeably. Although the semantic differences are subtle, privileged access management is the prevailing and the most current term (Haber, 2020, p. 91). Account and user management terms refer more to the history of PAM systems, which started from password management and extended via account management to access management (Carson, 2019).

Privileged access management is a part of the identity governance (IG) framework and complements IAM. Optimally, PAM is part of the organization's IAM model, but it can also be implemented independently. Although IAM and PAM provide logically same sort of functionality, PAM typically has its dedicated solutions and policies. The distinction between the two systems is that

PAM focuses on privileged access, typically meaning administrative use. In addition to the logical connection between IAM and PAM, IAM usually interfaces with PAM by managing the identities related to PAM accounts (Haber & Rolls, 2020, pp. 137-139).

Functionally privileged access management is a collection of methods to secure, manage, and monitor privileged access to resources. The components potentially found in a complete enterprise PAM system are shown in Figure 8, together with the accompanying features. PAM should not set any restrictions on the type of managed resources. Typical resource types are operating systems, applications, network devices, and databases. PAM ensures that only authorized users can administer secured systems and records the use of privileged accounts. PAM provides effective controls against external and internal threat actors. Internal actors gain legitimate access to systems with PAM, but the unalterable audit trail works as a deterrent for any deliberate misuse of privilege (Haber & Rolls, 2020).

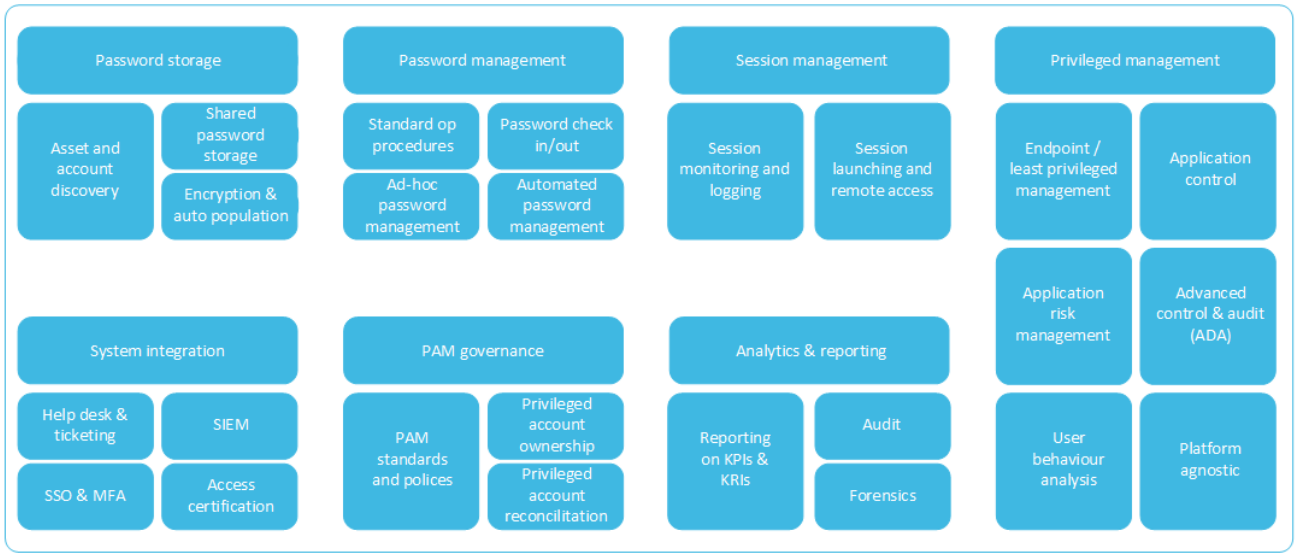

Figure 8 Components of a privileged access management system (adapted from Haber & Rolls, 2020, p. 140)

Session management provides the solutions to audit user interactions and creates indexable audit trails with the following features (Haber & Rolls, 2020, p. 143):

- Session monitoring and logging: Documents session activity to human-readable format for manual auditing and machine-readable format for log consolidation and further analytics such as event correlation or detecting indicators of compromise (IoC).
- Session launching: Launches a session and possibly injects credentials automatically.

Privileged management monitors, controls, and terminates privileged access to resources with the following features (Haber & Rolls, 2020, pp. 143-145):

- Endpoint least privilege management: Enforces least privilege model on targets and elevates privileges when needed.
- Application risk management: Assesses application risks before enabling privileged access.
- User behavior analysis (UBA): Detects suspicious usage patterns based on preset rules and learned baselines.
- Application control: Execution control based on whitelists, blacklists, and other quantifiable criteria.
- Advanced control and audit (ACA): Command filtering to prevent malicious activity. Can also block activity obfuscated from session monitoring, such as launching malicious scripts or subprocesses, using agent or client-based technology.
- Platform-agnostic: Protection for all platforms that use privileged access.

System integration provides the capability to integrate PAM with other security tools using the following features (Haber & Rolls, 2020, pp. 145-146):

- Helpdesk and ticketing: Integration to IT service management (ITSM) system for tracking and documenting the processes of granting or denying access.
- Single sign-on (SSO) and multifactor authentication (MFA): Prevents unauthorized access and mitigates various attack vectors such as password reuse and credential stuffing.
- Security information and event management (SIEM): Integration to centralized logging for advanced analytics and separation of duties.
- Access certification: Provides certification report on what was accessed and what was done.

Privileged access management governance manages the policies and procedures that an organization needs for its day-to-day operations using the following features (Haber & Rolls, 2020, pp. 146-147):

- PAM standards and policies: Govern who should have access to what and when. Uses role-based access control for granularity and can refer to role assignments in IAM.
- Privileged account ownership: Documents account ownership.
- Privileged account reconciliation: Keeps track of all aspects of PAM and identifies anomalies such as obsolete rules or accounts.

Analytics and reporting generates granular audit data with the following features (Haber & Rolls, 2020, p. 147):

- Reporting on key performance indicators (KPIs) and key risk indicators (KRIs): Continuous reporting on the health of the environment.
- Audit: Real time and historical reporting on all PAM solution use.
- Forensics: Detailed description on recorded events and granular reporting for investigating indicators of compromise and forensic investigation.

The number of features listed above is rather large and not all products have all the features. Similarly, not all organizations need all the listed features. It is important to form a tailored privileged

access strategy that implements only the controls that are relevant to the organization and its assets. Typically, the number of required controls increases with the size of the organization (Haber, 2020, p. 91-92).

The mere number of PAM features listed above suggests that deployment might be a complex task and it needs to be carefully planned to minimize negative effects to users. The benefits of a successful implementation certainly balance the effort needed for implementation. The motivation to implement a PAM solution can be approached using the main challenges that a PAM system can solve (Haber. 2020. pp. 93-97):

- Lack of a privileged account inventory. Decentralized and manual administration processes easily lead to forgotten and lost accounts with no control over their security.
- Lack of privileged account auditability. Even in the situation that every account is secured properly, and their use is within the guidelines, the organization cannot prove its compliance.
- Lack of visibility into shared accounts. Shared accounts are common in system administration, and without third-party tools it may be impossible to trace the users of such accounts.
- Lack of SSH key management. Unmanaged keys are prone to sharing and reuse, and rotating keys without proper tools is an administrative burden.

As a sign of need and interest for PAM functionality, multiple vendors provide their solutions. Gartner's privileged access management product review lists 41 different PAM vendors that have received reviews in the last 12 months (Gartner, 2022).

4 Description and analysis of the current PAM model

4.1 Network environment

To provide better understanding of the reasoning behind the current access management model and the organizational requirements, a brief introduction into the managed service provider network environment is in order. In general, the network structure of an MSP company resembles that of any other enterprise. Access from internet is strictly limited to only the necessary protocols, user endpoint health is monitored, and users must use strong authentication methods to get access. Access from internal networks to internet is also limited and protected with content filters. Internal networks are segmented according to functionalities and security levels and isolated from each other using firewalls.

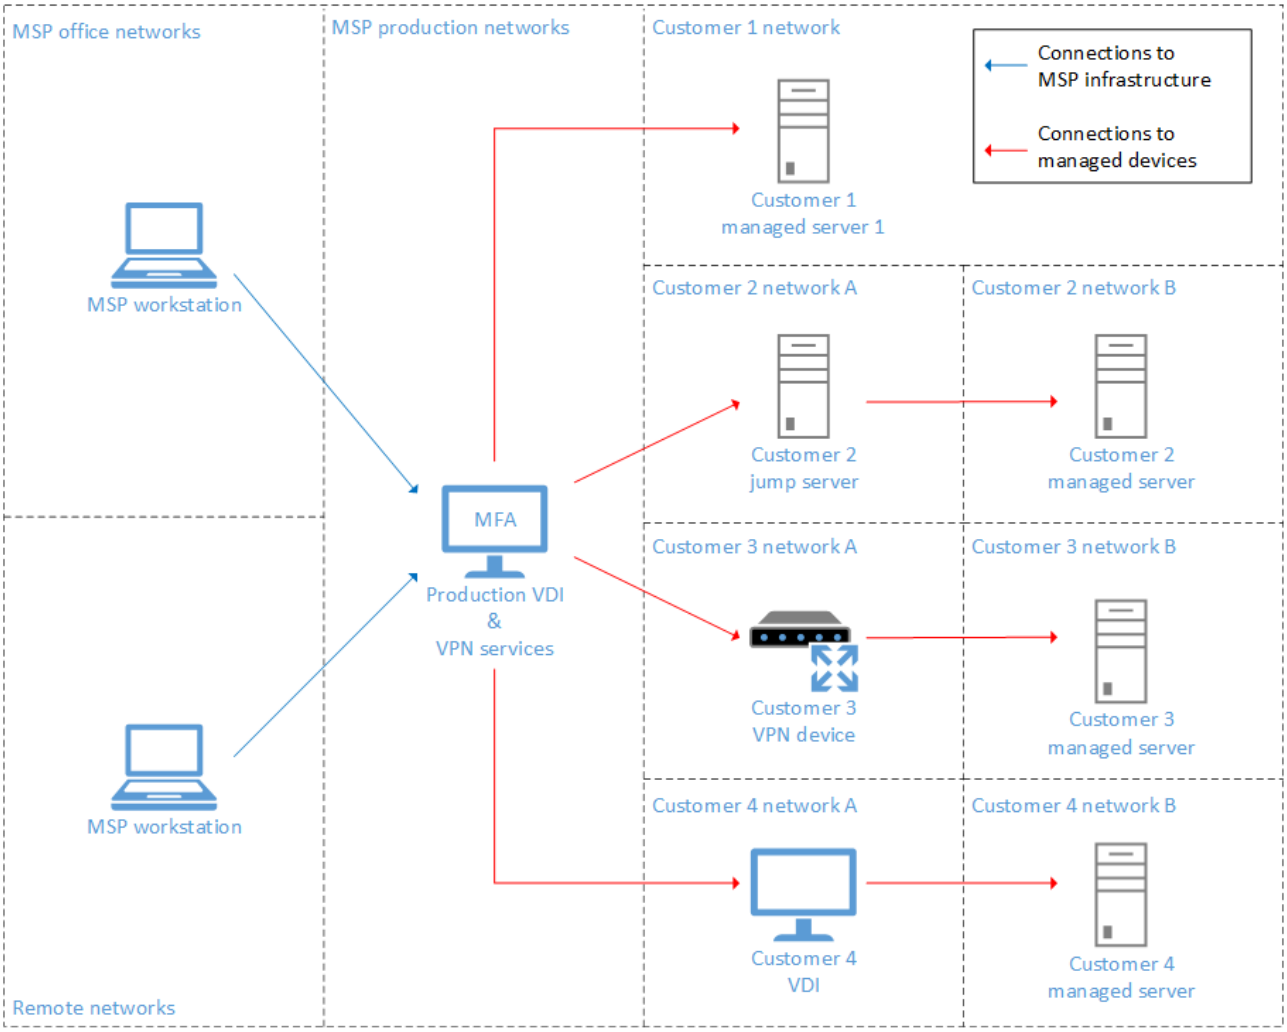

Figure 9 A simplified MSP network diagram

What separates MSPs from other enterprises is their access to the networks of other companies. Figure 9 displays a simplified diagram of MSP networks, their relation to the networks of managed customers, and connection paths for management connections in some typical connection scenarios. Office networks denote non-production networks inside the organization perimeter. One notable omission in the figure are the system-to-system connections from the MSP infrastructure to customers' infrastructure. Such connections are important enablers for system integration and automation, but out of scope of this study.

As illustrated in Figure 9, the connection techniques between MSP companies and their customers vary, but at least some sort of logical management connection must be established from the MSP network to the customer network. Jump servers, VPN connections, and virtual desktop infrastructure (VDI) are common ways to provide access to service providers.

Directly routed connectivity from MSP production networks to customer networks is a simple solution with potential for creating overlapping IP address spaces and routing issues as the number of customers and networks increases. A dedicated and isolated management network prevents such issues and provides a scalable and secure solution. A dedicated management network also improves availability by allowing access to the resources even when the customer networks are down. The disadvantage is that some managed devices do not allow adding a dedicated management interface and the management traffic must use a shared interface instead.

Jump servers or tool servers are placed in MSP managed networks that are reachable from MSP production network and have connectivity to a customer networks. Dedicated jump servers are good for restricting access but very poor in terms of scalability and administrative overhead. A shared jump server with access to several customers is a more scalable solution if permissions can be limited so that administrators only have access to the customers they are allowed to manage. VDI solutions are like jump servers regarding network configurations but do offer the advantage of integrated centralized management and improved scalability.

A point-to-site VPN connection is a traditional solution that still has its use in scenarios where other management network connections cannot be implemented. VPN connectivity has a high administrative overhead and multiple VPN connections and credentials are also a burden for the users.

Another technical factor that sets service providers apart from standard companies, is the need for credentials to access customer environments. Depending on the use case and customer requirements, the credentials can be personal or shared and either directory-based or local. Regardless of the credential type, they should be stored securely, and their usage should be audited and logged. As in the case of the management networks, using a different authentication source than the customer decouples the service provider from the customer infrastructure and provides better availability.

4.2 Organizational requirements

4.2.1 Documentation review

As the first part of analyzing the currently deployed model, organizational requirements are derived from relevant documentation. For most companies, the company security policy is a good starting point as it defines the basic stance of the enterprise cyber security. Even the public documentation usually outlines the high-level goals for security, and the principles that are used to achieve those goals. Roles and responsibilities are also described at least on the management level.

Selecting the public policies as the starting point serves two goals for the structure of this study. It helps to keep the work more general as the public enterprise policies probably vary less than internal policies. The other benefit is that the use of public sources does not disclose any information that might put the commissioner, Telia Cygate, at risk.

Since Telia Cygate is a subsidiary of Telia Company, it follows the enterprise level policies and guidelines of Telia Company. For Telia Company, the public security policy does not have any specific considerations for PAM, but it does have some general pointers that are relevant for the scope of this study (Telia, 2019). Perhaps the closest thing to a technical requirement, is the policy to provide business continuity plans for all business-critical systems. To describe the other policies in the document briefly, the company must seek to protect its data as well as the data of its customers, and to provide the agreed services to its customers in a secure manner.

Another public general security document of Telia Company, Security in Telia Company – general description, complements the public security policy and includes more specific instructions (Telia, 2020). The document describes how to use the well-known CIA triad of confidentiality, integrity, and availability to define the protection level of IT solutions. The general description also states that the company implements an information security management system (ISMS) based on the standard ISO/IEC 27001:2013 (International Organization for Standardization, 2013). The standard is a widely accepted framework for ISMS and covers policies and procedures for technical, physical, and legal controls. It is also noted that the ISMS has a statement of applicability (SoA) including all the security controls in ISO/IEC 27001:2013 Annex A. The Annex is a quite comprehensive

list of controls and provides a good basis for a check list that is needed to evaluate access management systems and controls.

The Security in Telia Company document also includes a chapter dedicated to access control. In general, access to information and assets must be controlled based on business, legal, and contractual factors. Other key points in the chapter are that access rights are only granted as necessary for an employee's responsibilities and duties, and they are assigned to personal user accounts. Access rights must be reviewed periodically or when the roles or duties of an employee change.

To dig deeper into organizational security rules and practices, one usually needs access to internal documents, as specific practices are often considered confidential. For this document review internal documentation in the company intranet was reviewed. Most of the internal guidelines and policies are the same or directly derived from the enterprise policies with one important addition. An internal document on information security architecture states that production system information security level is evaluated using CIS/SANS Top 20 (Telia Cygate, 2022). SANS Top 20 is the predecessor of CIS controls (SANS Institute, 2021). CIS controls, now in version 8, is a well-known framework and offers a comprehensive checklist of controls for creating cyber security defenses (Center for Internet Security, 2021).

4.2.2 Use case survey

The documentation review above establishes a rather general guidance on security policies that must be met when designing a privileged access management (PAM) system. To complement the documented policies with functional requirements, a use case survey is conducted. For this collection of requirements, a high-level use case analysis is sufficient. The main considerations are to find all the relevant actors and to list their use cases for the system. Figure 10 illustrates the simplified concept of relations between an actor and the system to be used as a basis for analysis. In this case, the users' motive to use the PAM system is the need to connect to resources to perform their job-related tasks. Here one needs to remember that the scope of the study omits machine-to-machine interactions and focuses solely on user actions.

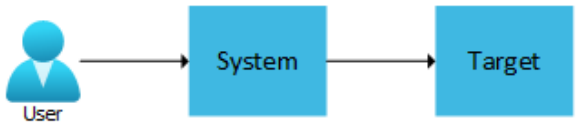

Figure 10 Simplified system function of a PAM solution

Since the scope of privileged access limits the target stakeholders to system administrators, and there is an existing system in place, the first step of identifying actors is straightforward. The system administrators can be divided into two subgroups, internal and external users. Technically, both groups share the same targets, goals, and use cases, but the administrative processes for managing external users are often different and needs to be considered. From the viewpoint of PAM system, external users might need some specific roles and permissions defined, but the core functionality requirements for the system are the same for internal and external users.

To describe a most generic use case for the system, it can be expressed in story format: “As a system administrator I need to connect to a resource to perform an administrative task”. The hypothesis is that all the use cases can be extrapolated from the generic one by enumerating all the different targets and tasks and placing them in the generic form to express more specific use cases.

The target resources are servers, network devices, and web services. To categorize targets in terms pertinent for the PAM system, network protocols used for management access are used here. The reasoning is that since the implemented PAM system only supports a limited number of protocols, it is easy to detect, which use cases require extra configuration to facilitate the use of unsupported protocols. Server management uses mostly RDP and SSH protocols and network device management uses mostly SSH and HTTP/S protocols, but telnet and TFTP are used as well. Web services use mostly HTTP/S protocol. Although RDP, SSH, and HTTP/S are ubiquitous and cover most user needs, it is important to include targets that use other protocols in the stories. Network device management with a thick client using vendor proprietary protocols or TFTP for file transfer and server or application management using Citrix ICA protocol are some examples of such user needs.

The network connectivity between the PAM system and the target is another factor separating use cases. For some use cases, connectivity from PAM system exists using standard routing. For some use cases there is no routable connection and management tasks are done using one or more

jump servers, VPN connections, or HTTPS tunneling. The PAM system supports the protocols listed above only for connections that the PAM servers can connect directly. An extra component can be deployed to remote networks that have connectivity to PAM and the remote targets to establish an encrypted communication channel. The extra component only supports RDP and SSH protocols currently, and therefore it is important to categorize the use cases accordingly.

Although there is a vast number of different administrative tasks that system administrators take care of daily, the pertinent tasks for the PAM system boils down to very few network level interactions. Again, it suffices to identify the access methods that are needed for accomplishing the tasks. The key task is to gain management access to a system using one of the protocols defined in the last chapter. As that access is provided, the user can perform various administration tasks, including copying textual data in and out using clipboard. Another key task is bidirectional file transfer between the user endpoint and the target. Depending on the management network protocol, this might be included in the management access or needs to use another channel.

Table 1 List of all detected use case parameters

ID	Target	Protocol	Network	Task
1	Server	RDP	Routable	Management access
2	Server	RDP	Routable	File transfer in/out
3	Server or device	SSH	Routable	Management access
4	Server or device	SSH	Routable	File transfer in/out
5	Service	HTTP/S	Routable	Management access
6	Service	HTTP/S	Routable	File transfer in/out
7	Server, device or service	Other	Routable	Management access
8	Server, device or service	Other	Routable	File transfer in/out
9	Server	RDP	Remote	Management access
10	Server	RDP	Remote	File transfer in/out
11	Server or device	SSH	Remote	Management access
12	Server or device	SSH	Remote	File transfer in/out
13	Service	HTTP/S	Remote	Management access
14	Service	HTTP/S	Remote	File transfer in/out
15	Server, device or service	Other	Remote	Management access
16	Server, device or service	Other	Remote	File transfer in/out

Table 1 lists all the parameters discussed above producing 16 different combinations. The number of different use cases sounds rather low at this point but it is imperative to remember that the combinations with the protocol other are a “catch-all” definition, that can generate considerably

larger number of more specific user stories. However, for the purposes of this requirement analysis, this level of detail is sufficient for the case analysis.

The use case list above does not consider one key feature of the selected PAM product, the capability to use either a web interface or native client programs for the supported protocols. Incorporating that distinction into Table 1 would double the number of use cases without adding much tangible value to the analysis. User stories are one suitable tool for describing the use of the system to such detailed level. To analyze the technical viability of the proposed solution, this study does not require the level of detail delivered by user stories. Yet, the use of web interface or native clients is discussed in the assessment since it does have some technical implications.

4.2.3 Check-list for requirements

Completing a checklist of requirements from different sources and pertaining to a single part of a large, complicated system that is responsible for organization-wide access management is no trivial task. In this chapter a checklist for PAM is compiled, bearing in mind that the access management platform includes network connectivity and connection filtering, identity management, several authentication providers, several password managers, and the PAM system. The services provided by all these systems intertwine in many places. For instance, password management is done in user directories, in a separate IAM solution, and in the PAM system. Some criteria in the resulting checklist might thus seem ambiguous, but they have been included if the PAM system is at least partially responsible for producing that control or service.

As noted in section 4.2.1, ISO/IEC 27001:2013 is used as a basis for building a checklist for the organizational requirements of PAM implementation. Similar approach in a larger scope is used by Purba & Soetomo (2018) in the context of establishing a baseline for choosing a suitable PAM product. This study is scoped more concisely and accordingly the checklist becomes more compact. ISO controls are complemented with CIS controls, as explained in section 4.2.1. The two frameworks have many overlapping controls as indicated in the CIS/ISO mapping (Center for Internet Security, 2023). The overlapping CIS controls are therefore omitted from the resulting checklist.

The resulting checklist is displayed in Appendix 1, where requirements sourced from ISO/IEC 27001:2013 use identifiers aligned with the source framework such as “A.9.1.2”. Requirements sourced from CIS controls are denoted by identifiers starting with C and a number referencing the source control number.

4.3 The current PAM model

4.3.1 Description

In the following presentation, company internal networks are categorized as office and production. Office networks host the systems and services that are used for daily office tasks such as employee workstations, printers, and internal services. Production networks host the systems and services that are used for providing the services that are consumed by customers and the internal users, office systems included. Of course, for some systems this leads to duality in categorization. For instance, a user can access the company email service with standard user account to read her or his emails, and on a different occasion, access the same system using a privileged account to perform system administration.

Currently implemented solution for access management is presented on a high-level in Figure 11. The figure includes the privileged access management of production infrastructure, customer servers, and customer network devices. The blue arrows represent user connections to the company infrastructure. The green arrows represent management connections audited with the PAM solution and the red arrows represent connections that are audited using other means. Network structure is displayed in a simplified form to keep the illustration readable, but the level of detail is sufficient to show the logical network locations of all the possible start points of management interactions. The figure also covers the use cases where users need to access resources only in office systems and are not affected by PAM. The figure omits some legacy access methods for the sake of readability. Third party access is also omitted since it follows the same path as the internal user access via virtual desktop infrastructure with the exception that external users have their dedicated VDI solution.

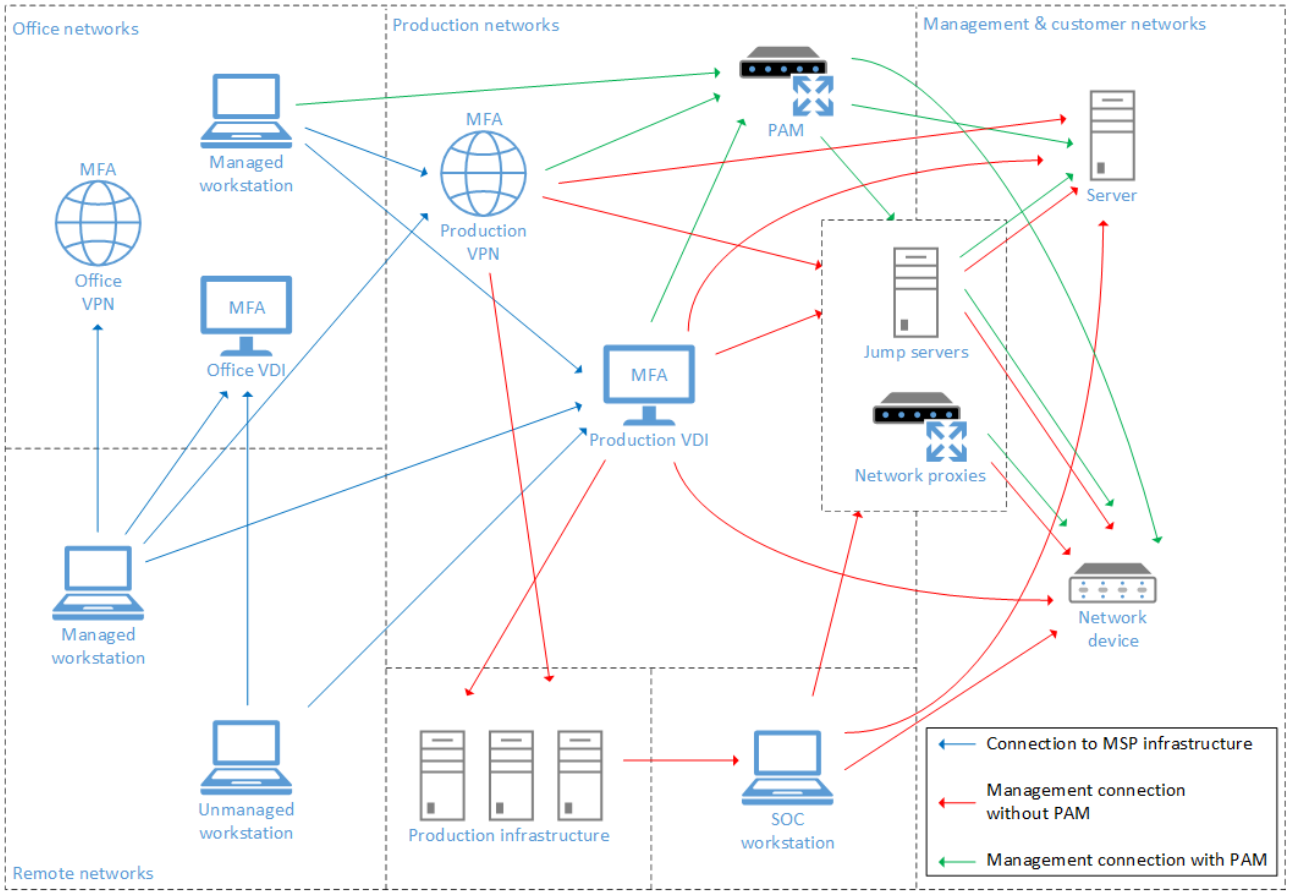

Figure 11 Current management model

The illustration covers how a user can access company resources remotely or from office networks. Remote connectivity is achieved using VDI or VPN. Only managed workstations can use VPN connections since the VPN solution must be able to verify the endpoint health information such as endpoint protection status. Both VPN and VDI have different implementations for office and production access to decouple the networks. VDI and VPN connections require strong authentication and are protected with multifactor authentication.

Identity management and provisioning of accounts, roles, and permissions is done mostly with an identity management product referenced here as IAM. Due to historical reasons, such as past mergers, and different function of network segments, there are multiple Active Directories in use for authentication. Non-windows systems use AD instances via LDAP and RADIUS protocols. Historical reasons also dictate that credential management is distributed to several tools potentially confusing the users.

The target systems of privileged access management mostly reside in the production and customer networks. Access to production infrastructure must always traverse via production VDI or VPN endpoints, even when the user is physically connected in the office network. As an exception, security operations center (SOC) has direct access from their workstations since their work area is physically separated from normal office space and has heightened physical security controls.

The main features of the PAM software are:

- support for RDP, SSH, VNC, and HTTP/S protocols using a web interface
- support for RDP and SSH native clients
- support for RDP and SSH protocol targets in remote networks using an extender component deployed in the target network
- capability to create video recordings of RDP, SSH, VNC, and HTTP/S sessions
- capability to include file transfers and clipboard use in audit trails
- granular role-based access control for target hosts
- support for passwords and SSH keys for local authentication to PAM
- support for passwords and Kerberos tokens for LDAP authentication to PAM
- support for storing target system credentials without exposing them to users
- support for rotating passwords in target systems
- access to all system settings using application programing interface (API)

The PAM system is deployed so that it is available from production networks only. PAM is configured to create a comprehensive audit trail of all connections and to store the trail files encrypted in a highly available external storage system. Connection targets are imported into the system from configuration management database (CMDB) via automation, but the system allows adding targets manually too.

The PAM product uses a role-based access management model for assigning permissions. Access roles are designed to produce a minimal number of groups while still respecting the principle of least privilege. A basic role only gives access to the PAM portal and all connections require additional roles. The additional roles fulfill the granularity requirements for internal production system management and different customer requirements. If a customer or a system requires a separate security clearance for access, the dependency is implemented in IAM. The user is granted access only after both the security clearance and the PAM role have been accepted in the IAM.

Considering the implementation of PAM, Figure 11 reveals ample room for improvement. PAM deployment is functionally complete, and policies are in place to direct users to use PAM when accessing customer servers. However, only a part of server access is enforced to PAM with firewall policies. As the number of arrows suggests, there are multiple ways to connect to managed

systems. As the company policy states that all access logs must be shipped to the centralized logging system, an audit trail is always created but the advanced capabilities of PAM are missed. For users the large number of ways to connect appear confusing and increase the need for documentation.

Customer network device management is mostly done via network proxies with SSH tunneling or by using proxies as jump servers. Proxied connections bypass PAM meaning that audit trail must be generated using other tools. The number of managed network devices is tens of thousands, so any changes to this management scenario merit careful consideration.

Access to the company's own production infrastructure is another connection scenario where PAM use is not always enforced. One reason for this is that the infrastructure needed to be installed before the PAM system could be implemented on top. Applying a restrictive control afterwards takes great care and involves a high risk of disruptions. Again, audit trail creation must take place outside the PAM system, increasing administrative overhead. From system administrator viewpoint, a centralized access management would streamline daily tasks as customer and internal systems could be managed coherently using a single point of access.

4.3.2 Compliance review

Since the use case list developed in section 4.2.2 and presented in Table 1 is partly based on the current management model, it is obvious that the current model does support all the listed use cases. Organizational requirements are reviewed using the checklist that was compiled in section 4.2.3 and is presented in Appendix 1. Assessing compliance with the checklist requires some interpretation for two reasons. Firstly, the checklist items are derived from cyber security frameworks and are therefore somewhat general in nature. Another reason is that although the checklist items are dependent of the PAM solution, they are also dependent on other systems in the management model. Therefore, the interpretation here is to check whether the current PAM solution enables compliance with the requirements given that all other components in management solution are configured accordingly.

Appendix 2 displays the results of organizational requirement review for the current PAM implementation. Compliant means that the requirement is fulfilled without notable exceptions or

shortcomings. Partially compliant means that compliance is possible but has some shortcomings. One example of such situation is a scenario, where controls are stated by policies and guidance, but not enforced technically. Noncompliant means that the control or its implementation is missing.

Out of 25 reviewed factors, six are compliant, 17 are partially compliant, and only two are deemed noncompliant. This is a satisfactory result given the fact that the PAM solution is retrofitted into a complex production environment and is fully implemented only for a subset of management scenarios. Still, many partial compliances leave plenty of room for improvement.

5 The proposed PAM model

5.1 Description

The aim is to create an improved access management model for privileged access. Of course, measuring the improvement might be rather ambiguous since usually a higher level of security means a lower level of usability. In this case the proposed changes potentially affect daily workflows of the administrators profoundly, so finding the right balance between security and usability is the key directive.

The existing management environment and the PAM deployment create a strong foundation for the development of an improved model. The drawback of a pre-existing environment that is already in production use is that implementing changes becomes more convoluted. The most important principles of this development are:

1. Adherence to company security requirements.
2. Compliance with the detected use cases.
3. Simplicity: Low administrative overhead for systems and documentation, ease of use for system administrators.
4. Modularity: Any products used should be changeable and vendor lock-in should be avoided.
5. Minimized administrative overhead: The solution must be robust and easy to administer.

Naturally, the goal of any PAM project must be that all access to protected assets uses PAM. This case is not an exception as PAM use is clearly mandated by company policies and requirements for auditing and access management. Comparison of the detected use cases and the features of the PAM product reveal that there should be no technical reasons to leave some use case scenarios

out. Most system administration work is done using the supported RDP, SSH, and HTTP/S protocols and any unsupported protocol can be included by using a jump server that is accessed via PAM.

The extra jump server is something that the users do like to avoid to make their jobs easier. In this case the extra work needed should not be significant since the PAM provides user-friendly passthrough authentication to servers in the same AD domain that is used for authenticating to the PAM itself. Furthermore, using a dedicated management server fulfills the concept of a privileged access workstation (PAW) described in section 3.4 and thus improves security. The jump servers do pose a challenge for audit trail searchability, though. Direct connections to targets are easy to search in the audit log using the target's name or IP address, but a connection via a jump server only shows the connection to the jump server searchable in the logs. Any connections from the jump server onward must be viewed from the video recording of the session, which does not have a text-based search.

Figure 12 represents the high-level design for the improved management model and PAM implementation. The aim for simplicity is evident when the number of connection arrows is compared to that of Figure 11. The main idea is to enforce the routing of every user-initiated management connection through the PAM system. The primary benefit is that all access control to the managed resources and audit log creation of management connections become unified and standardized. Centralized access control means better visibility into who can access what. All management connections and transactions including file transfers are recorded providing comprehensive audit trails. Log consolidation also makes searching the management connection logs easier, as all the logs are accessible from a single interface.

The secondary benefit is that this solution offers a single point of access for all administrative work instead of different customers having different methods for access. Administrative work in complex MSP environments often starts with confusion over how to get access to the target assets. The single point of access means that the administrator always knows the first step of accessing any managed asset. As all the PAM privileges are granted using an IAM integration, the process of requesting access is also clear. The unified management process should make documenting management connections simpler as well.

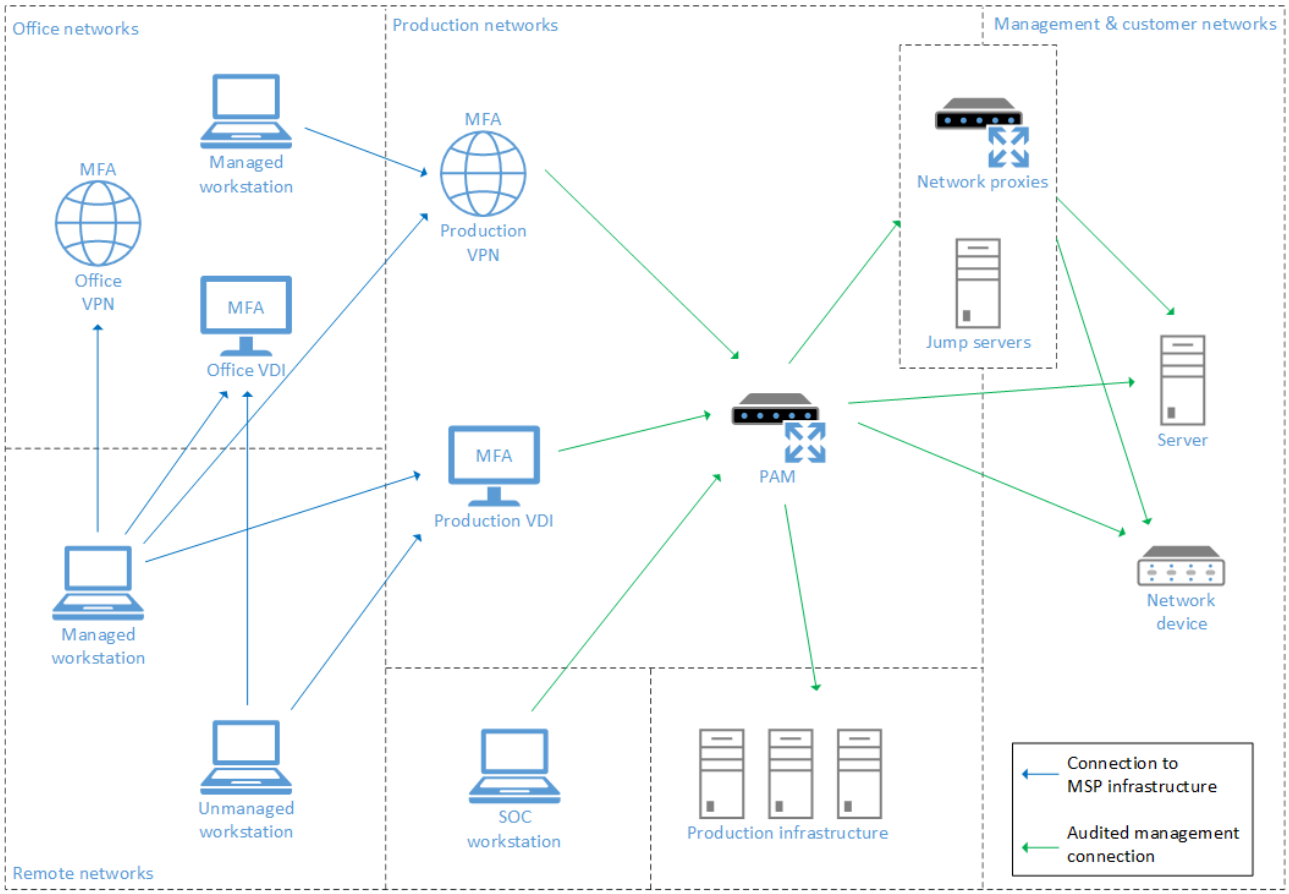

Figure 12 The proposed management model

As stated above, the proposed solution enforces practically all management connections to rely on privileges granted using the IAM solution. This brings some additional benefits as IAM provides account life cycle management, improved visibility into granted user permissions, and automated provisioning of permissions. User experience benefits too, as legacy domains outside IAM control are dropped and users need to manage fewer credentials. IAM also provides self-service password reset and an easy way to check users' own permissions.

Strong authentication is an important organizational requirement and MFA is always enforced at some point when accessing production networks. VPN and VDI connections both use an on-premise MFA provider. The PAM software does support MFA, but as there is no way accessing PAM directly from non-production networks, the MFA requirement is not enabled. Enabling MFA would improve security and be a step towards a zero trust architecture, described in section 3.4. It would also make the use of the system more tedious. If a less secure way of accessing PAM is implemented or security level needs to be raised, enabling MFA is quite simple.

The PAM product supports just-in-time access, which is a tenet of zero trust. Again, shifting to zero trust is a tempting goal, but this design does not use just-in-time access because of its reliance to ephemeral certificates. Issuing certificates requires a public key infrastructure, which is not necessarily available. Even for those domains that have PKI, the PAM must be able to connect to the domain, which is often complex due to the management network structure described in section 4.1.

5.2 Compliance review

When reviewed using the use case list developed in section 4.2.2 and presented in Table 1, the novel management model supports all the use cases same as the current model. Use cases 7, 8, 15, and 16 require the use of an extra jump server or VDI, but as explained in section 5.1., this should not be a significant issue.

As discussed at the end of section 4.2.2, the PAM product supports two ways of connecting: a web interface for all supported protocols and the use of native clients for RDP and SSH targets. The feature list in the two ways of connecting is mostly identical but the user experience is entirely different. Web client is clearly preferred by the product supplier, but many sysadmins still prefer using the native clients they are used to. SSH tunneling is one feature that is not supported in the web client but might be useful especially in management of network devices. From the auditing point of view, avoiding the use of tunneling is only a good thing, since data transferred in the tunnel cannot be audited.

Another important factor to consider is the feature set of the PAM extender component needed to support targets in remote networks. Currently the extender component only supports connection recording with the web client, which strongly negates the use of native clients for remote network targets. Furthermore, the extender component only supports RDP and SSH protocols, which means that the use of HTTP/S protocols needs to be done using a jump server or SSH tunneling.

In conclusion, the improved model supports all the detected use cases, but some cases require extra consideration and infrastructure. The non-native methods also require thorough testing with the stakeholders to find a way of work that creates an acceptable audit trail and does not affect productivity negatively.

The compliance with the organizational requirements is evaluated analogously to section 4.3.2, using the checklist compiled in section 4.2.3 and presented in Appendix 1. The result is displayed in Appendix 3. Compliance can be achieved in 21 out of 25 factors, and in the remaining four at least partial compliance is achievable.

Some of the improvements are not due to features of the PAM product but are enabled by better use of the PAM solution. One example of such a condition is the credential lifecycle management where the use of PAM implicitly means that the credentials must be controlled by the IAM system. IAM mitigates issues detected in checklist requirements A.9.2.2, A.9.2.4, and A.9.2.6.

Some factors are improved in the proposed model but cannot accomplish full compliance without changes in other systems in the management model. For instance, for requirements A.9.4.3 and C.5.2 the use of PAM will force the use of IAM and some password quality control for accounts that are used for authentication to PAM. However, PAM does not include any password quality control on the target credentials. Requirement A.9.2.3 is another example of such case. The PAM implementation supports division of regular and admin accounts and makes the separation easier as users login to the PAM system with domain user accounts and can then use privileged accounts for target logins. This improvement also requires that supporting systems and targets are configured accordingly. Requirement A.9.3.1 is a case where PAM brings added value as target credentials need not be exposed to users anymore. Yet, the authoritative credential storage should be in a password vault service.

In comparison to the existing solution, the novel approach improves compliance with the requirement checklist. In 16 out of 25 checklist items, the compliance status improves. The most notable improvements are related to the centralization of management traffic to one access point, that is the PAM solution. A single access point makes access management including revoking access more efficient, as measured with checklist items A.9.1.2, A.9.2.6, and C.6.7. Checklist items A12.4.3, C6.7, C8.5 and C8.8 are related to auditing, and they all benefit from the elimination of alternative management access paths as a single point of access is easier to monitor and provides a comprehensive audit trail of activity in a centralized view.

The conclusion of the organizational requirement review is that the improved model has better compliance to the requirements checklist than the present model. There are still some factors where a complete compliance cannot be reached, but mitigation and improvement methods are already discussed above.

5.3 Implementation

Most of the work needed for implementing the proposed change must be done on network level. This means modifying firewall rules to block access to managed resources if the source is not the PAM system. Since the blocking should affect only the privileged access of human users and not machine-to-machine communications, there are considerable risks involved. RDP traffic is usually sourced from users and is therefore safe to block, but SSH and HTTP/S are often used by automated system management and monitoring, and erroneous blocking could have serious consequences. Consequently, not being able to block all SSH traffic circumventing the PAM solution does leave a potential gap that must be audited using other methods.

Besides the network changes, there are other caveats to consider when designing the transition to the improved model. First and foremost, the transition means a significant change in some sysadmin workflows. Resistance of change is something that will need to be responded with timely and comprehensive communications and user training. Creating easily accessible documentation of system usage and its features must be an integral part of the user training.

As the use of PAM will be expanded to network devices, there must be a strategy and methods on how to register tens of thousands of devices in PAM without the need for users to type them in manually. Network devices are also typically in remote networks that are not directly accessible from the PAM. For this purpose, the PAM product has an extender component that can be deployed on a Linux server and creates an SSH tunnel to the PAM. Connections to networks accessible to the extender server can then be routed through that tunnel. Distribution of the extenders needs to be planned and executed in co-operation with the network admins.

The change is quite pervasive in the sense that involvement is required from many parts of the organization, including some that do not participate in the daily administration of customer assets. Many processes that involve importing or creating new managed resources need to be adjusted to

the change in the management model. The implications of the novel model also need to be communicated to people that need to make decisions regarding management access. One example are the network teams that oversee the firewall policy administration. They need to be informed of the policy that disallows management connections that are not audited with the PAM solution.

6 Conclusions

6.1 Results and reliability

The primary objective of this work was to design a more secure access management model. A special emphasis was on the role the privileged access management (PAM). The related research question was whether the novel model fulfills the organizational requirements. Discovering and understanding the organizational requirements is a substantiative task itself, and in this work, it is conducted with a documentation review leading to a custom checklist of requirements.

The requirement checklist is constructed in section 4.2.3 and the compliance review is conducted in section 5.2. Quantitatively assessing, only four out of 25 checklist items are left in partially compliant status instead of fully compliant. The result is that the proposed model fulfills the requirements. Another conclusion is that the PAM system is a central component in this solution. It would be very difficult to enable corresponding features without a PAM system, especially in the audit domain.

The secondary research question, “how the novel approach improves security”, is another way to evaluate the proposed model. The compliance comparison of novel and existing solutions is conducted in section 5.2. The result indicates improved security especially in access management and auditing functions. It is important to notice that these improvements are mainly induced by the enforcement of a single point of management access with less permissive firewall policies.

Another objective of the work was to document the most important design factors and considerations of the proposed model including the rationale of the decisions. A necessary extension of this objective was that the documentation should provide an introduction into identity and access management and privileged access management to help understand the key concepts, to provide suitable background information into the PAM design process, and to offer references for further

development. The proposed access management model design is described at length in section 5.2 including the arguments for the decisions. Chapter 3 provides the theoretical foundation for all the key concepts discussed in subsequent chapters. The list of references provides the sources for the theory and a good starting point for researching identity and access management further.

The qualitative review of current and improved PAM models shows that most compliance issues in the current model can be mitigated in the improved model. Generally, this result looks reliable and makes sense, as all privileged use is transitioned to use PAM and more strict security controls are applied. The procedure to arrive at the results and conclusions is described in detail and for most parts is derived from industry standards. Therefore, the work should be repeatable, consistent, and generalizable, fulfilling the requirements of valid research.

There is some uncertainty and room for discussion though. Since the organizational requirements are derived from frameworks that are quite generic, the custom checklist used for compliance review leaves some margin for interpretation. Another source for uncertainty is that the evaluation of compliance is done in the context of a large and complex access management system with strong interdependence in features. Arguably, the criteria could be selected to be more pertinent to the PAM only, but such an approach might not serve the purpose of securing access altogether.

6.2 Discussion and further development

The most obvious need for further development in implementing a more secure access management model is to further elaborate the transition process to include all managed targets in the PAM solution. As described in section 5.3, the pre-existing environment poses multiple challenges for the implementation. This underlines the importance of including a secure management model as a native part of any system design process from the begining.

The most interesting question relating to the improved access management model that was not discussed in the previous chapters is the relation and ratio of security and usability. Addressing the question was left to this final section as it is difficult to find measurable quantities that would help evaluate the issue.

The subject of this study affects the daily workflows of numerous professionals, and there is no doubt that not all of them embrace the novel approach as it introduces changes in very basic tasks such as connecting and authenticating to managed servers. The positive effect of implementing the more secure model is that all the management must be done using a single gateway. The current model has multiple ways and routes of initiating management connections, which is confusing and increases the need for documentation. In the novel model the starting point for management connections is always the same, which streamlines workflows and requires less documentation. This difference is evident when viewing Figures 11 and 12. Still, the limitations of the PAM features, as discussed in section 5.1, sometimes add extra twists to workflows especially when using other protocols than RDP and SSH or accessing targets in non-routable networks.

Another objective that needs further attention is the documentation of the access management solution. This work provides the high-level documentation including the reasoning for design decisions. The system integration and successful daily operation demand a more detailed and non-public system documentation that must be done internally within Telia Cygate, the commissioner of this work.

One topic to consider as a future project is how to better utilize the role-based access control in PAM. Now identity management is done in a separate IAM that provisions memberships to Active Directory security groups using high granularity. The approach serves well the principle of least privilege as the number of assigned but unneeded privileges is minimal. The drawback is that most users have tens of privileges assigned to them and determining one's overall privileges is difficult. New users also often need several iterations to get the correct permissions assigned. If the organizational structure allows it, it would be easier and more efficient to define roles with preset permission set at least for the most common administrative roles.

In addition to the RBAC functions, the PAM software includes many of the features described in section 3.6 that are not currently utilized. Section 5.1 explains the decisions to omit some of the features, but the PAM software has functionality that merits further evaluation. For example, password and key rotation and providing short-term access upon request offer interesting chances for improving security and administrative workflows. Evaluation and possible implementation of these features makes sense as they are part of the software license that has been paid for.

References

Bassett, G., Hylender, C. D., Langlois, P., Pinto, A. & Widup, S. (2022). *Data breach investigations report*. Verizon.com. https://www.verizon.com/business/resources/reports/2022/dbir/2022-data-breach-investigations-report-dbir.pdf

Bhattacherjee, A. (2012). *Social science research: Principles, methods, and practices*. Digitalcommons.usf.edu. https://digitalcommons.usf.edu/cgi/viewcontent.cgi?article=1002&context=oa_textbooks

Bkav. (2017, November 10). *Bkav's new mask beats Face ID in "twin way": Severity level raised, do not use Face ID in business transactions*. https://www.bkav.com/en/top-news/-/view-content/65202/bkav-s-new-mask-beats-face-id-in-twin-way-severity-level-raised-do-not-use-face-id-in-business-transactions

Bryman, A. (2012). *Social research methods (4th ed.)*. Oxford University Press.

Carson, J. (2019). *Password managers to privileged access management: The evolution*. Delinea blog. https://delinea.com/blog/privileged-access-vs-account-management

Center for Internet Security. (2021). *CIS controls version 8*. https://www.cisecurity.org/controls/v8/

Center for Internet Security. (2023). *CIS controls v8 Mapping to ISO/IEC 27001:2022*. https://www.cisecurity.org/insights/white-papers/cis-controls-v8-mapping-to-iso-iec-27001-2022

Chapple, M. (2021). *Access control and identity management (3rd ed.)*. Jones & Bartlett Learning.

Department of defense. (1985, December 26). *DoD 5200.28-STD - Department of defense trusted computer system evaluation criteria*. https://csrc.nist.gov/csrc/media/publications/conference-paper/1998/10/08/proceedings-of-the-21st-nissc-1998/documents/early-cs-papers/dod85.pdf

Dibbern, J., Goles, T., Hirschheim, R., & Jayatilaka, B. (2004). Information systems outsourcing: A survey and analysis of the literature. *Database for Advances in Information Systems*, 35(4), 6-102.

Frank. (2013, September 21). *Chaos computer club breaks Apple TouchID*. Chaos Computer Club. https://www.ccc.de/en/updates/2013/ccc-breaks-apple-touchid

FIDO Alliance. (2022, May 5). *Apple, Google and Microsoft commit to expanded support for FIDO standard to accelerate availability of passwordless sign-ins*. https://fidoalliance.org/apple-google-and-microsoft-commit-to-expanded-support-for-fido-standard-to-accelerate-availability-of-passwordless-sign-ins/

Gartner. (2023, January 3). *Identity management systems reviews*. Gartner Peer Insights. https://www.gartner.com/reviews/market/identity-governance-administration

Grand View Research. (2022). *IT services outsourcing market size, share & trends analysis report by service (application, emerging technology), by location (on-shore, off-shore), by end-use, and segment forecasts, 2020 – 2027*. https://www.grandviewresearch.com/industry-analysis/it-services-outsourcing-market

Grassi, P. A., Garcia, E. & Fenton, J. (2017). *NIST special publication 800-63-3 - Digital identity guidelines*. https://doi.org/10.6028/NIST.SP.800-63-3

Grassi, P. A., Fenton, J., Newton, E., Perlner, R., Regenscheid, A., Burr, W., Richer, J., Lefkovitz, N., Danker, J., Choong, Y-Y., Greene, K & Theofanos, M. (2017). *NIST special publication 800-63B - Digital identity guidelines - Authentication and lifecycle management*. https://doi.org/10.6028/NIST.SP.800-63b

Grassi, P. A., Richer, J. P., Squire, S. K., Fenton, J. L., Nadeau, E. M., Lefkovitz, N. B., Danker, J. M., Choong, Y., Greene, K. K., & Theofanos, M. F. (2017). *NIST special publication 800-63C - Digital identity guidelines - Federation and assertions*. https://doi.org/10.6028/NIST.SP.800-63c

Haber, M.J. (2020). *Privileged attack vectors*. Apress.

Haber, M.J. & Rolls, D. (2020). *Identity attack vectors: Implementing an effective identity and access management solution*. Apress.

Hopkins, N. (2017). *Deloitte hit by cyber-attack revealing clients' secret emails*. The Guardian. https://www.theguardian.com/business/2017/sep/25/deloitte-hit-by-cyber-attack-revealing-clients-secret-emails

International Organization for Standardization. (2013). *Information technology — Security techniques — Information security management systems — Requirements*. https://www.iso.org/standard/54534.html

International Telecommunication Union. (2009). *Recommendation ITU-T Y.2720*. https://www.itu.int/rec/T-REC-Y.2720-200901-I

Jacobson, I., Spence, I. & Bittner K. (2011). *Use-case 2.0. The guide to succeeding with use cases*. Ivar Jacobson International. https://www.ivarjacobson.com/files/field_iji_file/article/use-case_2_0_jan11.pdf

Jakkal, V. (2021). *The passwordless future is here for your Microsoft account*. Microsoft Security. https://www.microsoft.com/security/blog/2021/09/15/the-passwordless-future-is-here-for-your-microsoft-account/

Microsoft. (2014, November 19). *Active Directory structure and storage technologies*. Microsoft Learn. https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc759186(v=ws.10)

Microsoft. (2022a, August 17). *Active Directory domain services overview*. Microsoft Learn. https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/get-started/virtual-dc/active-directory-domain-services-overvie

Microsoft. (2022b, September 20) *Security principals*. Microsoft Learn. https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/manage/understand-security-principals

Microsoft. (2022c, December 1). *Segmentation strategies*. Microsoft Learn. https://learn.microsoft.com/en-us/azure/architecture/framework/security/design-segmentation

Microsoft. (2023, March 4). *Developing a privileged access strategy*. Microsoft Learn. https://learn.microsoft.com/en-us/security/privileged-access-workstations/privileged-access-strategy

MITRE. (2021, August 31). *Use alternate authentication material: Pass the hash*. https://attack.mitre.org/techniques/T1550/002/

Neuman, C., Yu, T., Hartman & S., Raeburn, K. (2005, July). *The Kerberos network authentication service (V5)*. IETF. https://www.ietf.org/rfc/rfc4120.txt

National Institute of Standards and Technology. (2020). *NIST special publication 800-53 revision 5. Security and privacy controls for information systems and organizations*. https://doi.org/10.6028/NIST.SP.800-53r5

OWASP. (2021). *Password storage cheat sheet*. OWASP cheat sheet series. https://cheatsheetseries.owasp.org/cheatsheets/Password_Storage_Cheat_Sheet.html

Prince, M., Stinson-Diess D. & Zaman S. (2022). *The mechanics of a sophisticated phishing scam and how we stopped it*. The Cloudflare blog. https://blog.cloudflare.com/2022-07-sms-phishing-attacks/

Purba, A. & Soetomo, M. (2018). Assessing privileged access management (PAM) using ISO 27001:2013 control. *ACMIT Proceedings*, 5, 65-76. http://dx.doi.org/10.33555/acmit.v5i1.76

Rose, S., Borchert, O., Mitchell, S. & Connelly S. (2020). *NIST special publication 800-207 - Zero trust architecture*. https://doi.org/10.6028/NIST.SP.800-207

Saltzer, J. & Schroeder, M. (1975). *The protection of information in computer systems*. University of Colorado Springs. http://cs.uccs.edu/~cs691/designPrinciples/ProtectionOfInformationInComputerSystems1975.pdf

SANS Institute. (2021, May 18). *CIS controls v8*. Blog. https://www.sans.org/blog/cis-controls-v8/

Saunders, M. N. K., Lewis, P. & Thornhill, A. (2019). *Research methods for business students (8th edition)*. Pearson.

Semersheim, J. (Ed.). (2006, June). *Lightweight directory access protocol (LDAP): The protocol*. IETF. https://datatracker.ietf.org/doc/html/rfc4511

Smilyanets, D. (2022, April 28). *Session hijacking and MFA bypass*. Recorded future. https://www.recordedfuture.com/session-hijacking-mfa-bypass

Telia. (2019). *Group policy – Security*. https://www.teliacompany.com/globalassets/telia-company/documents/about-telia-company/public-policy/group-policy---security.pdf

Telia. (2020). *Security in Telia company – General description*. https://www.teliacompany.com/globalassets/telia-company/documents/about-telia-company/what-we-do/security-in-telia-company.pdf

Telia Cygate. (2022). *Tietoturva-arkkitehtuuri [Information security architecture]*. Unpublished internal company document.

Telia Cygate. (2023). *Telia Cygate | Telia yrityksenä*. https://www.telia.fi/telia-yrityksena/telia-cygate

Toikko, T. & Rantanen, T. (2009). *Tutkimuksellinen kehittämistoiminta: Näkökulmia kehittämisprosessiin, osallistamiseen ja tiedontuotantoon [Research-based development: Viewpoints into development process, involvement, and knowledge production]*. Tampere University Press.

Vinberg, S. & Overson, J. (2021). *2021 Credential stuffing report*. F5.com. https://www.f5.com/content/dam/f5-labs-v2/article/pdfs/2021-Credential-Stuffing-Report-rev-28APR21.pdf

Ware, W. H. (1970). *Security controls for computer systems: Report of defense science board task force on computer cecurity*. Rand. https://www.rand.org/pubs/reports/R609-1.html

Windley, P. J. (2023). *Learning digital identity*. O'Reilly Media, Inc.

Appendices

Appendix 1. Organizational requirements checklist

Identifier	Requirement header	Objective
A.9.1.2	Access to networks and network services	Users shall only be provided with access to the network and network services that they have been specifically authorized to use.
A.9.2.2	User access provisioning	A formal user access provisioning process shall be implemented to assign or revoke access rights for all user types to all systems and services.
A.9.2.3	Management of privileged access rights	The allocation and use of privileged access rights shall be restricted, controlled, and use user ID's different from those used for regular business activities.
A.9.2.4	Management of secret authentication information of users	The allocation of secret authentication information shall be controlled through a formal management process.
A.9.2.6	Removal or adjustment of access rights	The access rights of all users to information and information processing facilities shall be removed upon termination of their employment or agreement or adjusted upon change.
A.9.3.1	Use of secret authentication information	Users shall be required to follow the organization's rules in the use of secret authentication information, e.g., handling of passwords.
A.9.4.1	Information access restriction	Access to information and application system functions shall be restricted in accordance with the access control rules.
A.9.4.2	Secure log-on procedures	Where required by the access control rules, access to systems and applications shall be controlled by a secure log-on procedure.
A.9.4.3	Password enforcement and management	Applications using passwords shall ensure quality and interactive management of passwords
A.10.1.2	Key management	Rules on the use, protection, and lifetime of cryptographic keys shall be developed and implemented through their whole lifecycle.
A.12.4.3	Administrator and operator logs	System administrator and system operator activities shall be logged and the logs protected and regularly reviewed
A.16.1.7	Collection of evidence	The organization shall define and apply procedures for the identification, collection, acquisition, and preservation of information, which can serve as evidence.
A.18.1.3	Protection of records	Records shall be protected from loss, destruction, falsification, unauthorized access, and unauthorized release, in accordance with legislative, regulatory, contractual, and business requirements.
C.4.6	Securely manage enterprise assets and software	Access administrative interfaces over secure network protocols, such as SSH/HTTPS. Do not use insecure protocols, such as Telnet and HTTP, unless operationally essential.
C.5.1	Establish and maintain an inventory of accounts	Establish and maintain an inventory of all accounts managed in the enterprise.

C.5.2	Use unique passwords	Use unique passwords for all enterprise assets.
C.5.3	Disable dormant accounts	Delete or disable any dormant accounts after a period of 45 days of inactivity.
C.5.4	Restrict administrator privileges to dedicated administrator accounts	Restrict administrator privileges to dedicated administrator accounts and use non-privileged accounts for other purposes.
C.5.6	Centralize account management	Centralize account management through a directory or identity service.
C.6.7	Centralize access control	Centralize access control for all enterprise assets through a directory service or SSO provider, where supported.
C.6.8	Define and maintain role-based access control	Define and maintain role-based access control, through determining and documenting the access rights necessary for each role to successfully carry out its assigned duties.
C.8.5	Collect detailed audit logs	Configure detailed audit logs for assets containing sensitive data. Include event source, date, username, timestamp, addresses, and other useful elements that could assist in a forensic investigation.
C.8.8	Collect command-line audit logs	Collect command-line audit logs. Example implementations include collecting audit logs from PowerShell, BASH, and remote administrative terminals.
C.8.9	Centralize audit logs	Centralize, to the extent possible, audit log collection and retention across enterprise assets.
C.8.12	Collect service provider logs	Collect service provider logs, if supported. Example implementations include collecting authentication and authorization events, data creation and disposal events, and user management events.

Appendix 2. Current PAM solution compliance

Identifier	Requirement header	Current PAM solution compliance
A.9.1.2	Access to networks and network services	Partially compliant
A.9.2.2	User access provisioning	Partially compliant
A.9.2.3	Management of privileged access rights	Partially compliant
A.9.2.4	Management of secret authentication information of users	Noncompliant
A.9.2.6	Removal or adjustment of access rights	Partially compliant
A.9.3.1	Use of secret authentication information	Partially compliant
A.9.4.1	Information access restriction	Compliant
A.9.4.2	Secure log-on procedures	Compliant
A.9.4.3	Password enforcement and management	Partially compliant
A.10.1.2	Key management	Noncompliant
A.12.4.3	Administrator and operator logs	Partially compliant
A.16.1.7	Collection of evidence	Partially compliant
A.18.1.3	Protection of records	Partially compliant
C.4.6	Securely manage enterprise assets and software	Compliant
C.5.1	Establish and maintain an inventory of accounts	Partially compliant
C.5.2	Use unique passwords	Partially compliant
C.5.3	Disable dormant accounts	Partially compliant
C.5.4	Restrict administrator privileges to dedicated administrator accounts	Partially compliant
C.5.6	Centralize account management	Compliant
C.6.7	Centralize access control	Partially compliant
C.6.8	Define and maintain role-based access control	Partially compliant
C.8.5	Collect detailed audit logs	Partially compliant
C.8.8	Collect command-line audit logs	Partially compliant
C.8.9	Centralize audit logs	Compliant
C.8.12	Collect service provider logs	Partially compliant

Appendix 3. Proposed PAM solution compliance

Identifier	Requirement header	Proposed PAM solution compliance
A.9.1.2	Access to networks and network services	Compliant
A.9.2.2	User access provisioning	Compliant
A.9.2.3	Management of privileged access rights	Partially compliant
A.9.2.4	Management of secret authentication information of users	Compliant
A.9.2.6	Removal or adjustment of access rights	Compliant
A.9.3.1	Use of secret authentication information	Partially compliant
A.9.4.1	Information access restriction	Compliant
A.9.4.2	Secure log-on procedures	Compliant
A.9.4.3	Password enforcement and management	Partially compliant
A.10.1.2	Key management	Compliant
A.12.4.3	Administrator and operator logs	Compliant
A.16.1.7	Collection of evidence	Compliant
A.18.1.3	Protection of records	Compliant
C.4.6	Securely manage enterprise assets and software	Compliant
C.5.1	Establish and maintain an inventory of accounts	Compliant
C.5.2	Use unique passwords	Partially compliant
C.5.3	Disable dormant accounts	Compliant
C.5.4	Restrict administrator privileges to dedicated administrator accounts	Compliant
C.5.6	Centralize account management	Compliant
C.6.7	Centralize access control	Compliant
C.6.8	Define and maintain role-based access control	Compliant
C.8.5	Collect detailed audit logs	Compliant
C.8.8	Collect command-line audit logs	Compliant
C.8.9	Centralize audit logs	Compliant
C.8.12	Collect service provider logs	Compliant