

Security requirement

SaaS Public Cloud Usage

Deutsche Telekom Group

Version	1.0
Date	Jul 1, 2025
Status	Released

Publication Details

Published by
Deutsche Telekom AG
Vorstandsbereich Technology & Innovation
Chief Security Officer

Reuterstrasse 65, 53315 Bonn
Germany

File name	Document number 3.82	Document type Security requirement
Version 1.0	State Jul 1, 2025	Status Released
Contact Telekom Security psa.telekom.de	Validity Jul 1, 2025 - Jun 30, 2030	Released by Stefan Pütz, Leiter SEC-T-TST

Summary
This document covers secure usage of SaaS public cloud services. The requirements in this document are mostly given on high level so they can cover most of SaaS cloud services providers and use cases, but the document does not provide technical implementation details for a specific cloud provider.

Copyright © 2025 by Deutsche Telekom AG.
All rights reserved.

Table of Contents

1.	Introduction	4
2.	Introduction to Cloud Computing	5
3.	Public Cloud Usage Prerequisites	7
4.	Governance and Management	12
4.1.	Session Protection	17
5.	Authentication and Authorization	19
6.	Authentication attribute "password"	30
7.	Protecting Data and Information	34
8.	Protecting Availability and Integrity	38
9.	Architecture and Networking	40
9.1.	Cloud Connectivity	41
9.2.	Cloud Connectivity - External Connectivity	41
10.	Logging and Monitoring	44
11.	References	50
12.	Glossary	51

1. Introduction

This security document has been prepared based on the general security policies of the Group.

The security requirement is used as a basis for an approval in the PSA process, among other things. It also serves as an implementation standard for units which do not participate in the PSA process. These requirements shall be taken into account from the very beginning, including during the planning and decision-making processes. When implementing these security requirements, the precedence of national, international and supranational law shall be observed.

If compliance with the described requirements can not be achieved or is only partially feasible in individual cases, risk assessments must be carried out together with a Security and/or Data Privacy Expert (in accordance with the relevant requirements) and possible alternative protective measures must be agreed.

2. Introduction to Cloud Computing

Introduction

According to NIST (National Institute of Standards and Technology), cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models, and four deployment models.

Scope

This document covers the following topics:

- security guidelines for usage of SaaS public cloud services, including APIs
- security guidelines for cloud management
- security guidelines for cloud users and their applications

for the following type of cloud environments:

- public cloud

The following cloud types are not in scope of this document:

- private clouds
- community clouds

Hybrid clouds are not covered in particular. If an application uses both public and private cloud, public cloud requirements apply to all parts of applications which are located in public cloud.

In scope:

- Public cloud API, usage, configuration, limits
- SaaS public cloud services

Out of scope of this document:

- IaaS and PaaS public cloud services
- Operating system and application requirements
- Container and orchestration technologies (like Docker and Kubernetes)

Shared Responsibility Model

Security of the cloud and services provided on top of cloud is job for both cloud service provider and cloud consumer. The security of the cloud platform and offered services is a responsibility of the cloud service provider, while their use, either direct or as building blocks for a final service is responsibility of the cloud consumer (called security in the cloud). The boundary depends on the service model used, e.g. when using IaaS, cloud provider needs to secure hypervisor, and separate network and storage between tenants, while cloud consumer needs to properly secure virtual machines and software running inside them.

This approach is called shared responsibility model and one visual example can be seen here <https://docs.microsoft.com/en-us/azure/security/fundamentals/shared-responsibility>.

It is worth noting that even with shared responsibility model, the cloud customer remains accountable for the data security.

Background Story

Definition of information security according to ISO/IEC 27000:2014 is: "Preservation of confidentiality, integrity and availability of information. Note: In addition, other properties, such as authenticity, accountability, non-repudiation and reliability can also be involved."

Threats, either deliberate or coming from environment, accidents, negligence are typically classified in the following types:

- physical damage

- natural events
- loss of essential services
- compromise of information
- technical failures
- compromise of functions

Protection mechanisms against threats may include protection mechanisms of the cloud service provider (which must be still configured accordingly by customers) and those which customer implements on top of other services provided by the cloud provider, or naturally - combination of both.

In a cloud environment, like in any other data processing environment, it is also necessary to classify the data and to consider not only data availability but also the data lifecycle (e.g. when the data can be archived, deleted etc.) and respectively protect the data during the whole lifecycle.

3. Public Cloud Usage Prerequisites

Cloud Service Provider (CSP) provides computer system resources (e.g. compute, storage, network and application) to customers. Resources can be provisioned through APIs by customer. CSP is responsible for security of their infrastructure, services and APIs, while customer bears responsibility to secure everything built on top of resources and APIs provided by CSP.

Req 1	Before starting to work in public cloud, all users of cloud services must examine the relevant documentation.
-------	---

All cloud service users, including maintainers of the landing zone and especially application/system maintainers and users, must make themselves familiar with the cloud services which they are making available or using. Landing zone operators may also prepare relevant information for users of the public cloud to facilitate easier onboarding (or this can be made a formality - as part of acceptance of terms and conditions for using the public cloud).

Motivation: The security of cloud services of the public cloud service providers is generally good because these services are public and under constant attack, and most security incidents are not due to the vulnerabilities in the services, but due to the wrong configuration, misuse or human errors. Some typical examples include:

- *(inadvertently) exposing private object storage bucket to public*
- *creation of allow-all security groups (inbound and outbound), e.g. exposing internal systems (e.g. databases) to Internet*
- *(inadvertently) exposing credentials of cloud services (enabled debug output, keys in plain text in CI/CD chain)*

Implementation example: Relevant documentation includes, but is not limited to:

- approval of a specific CSP and the cloud service inside DTAG, including the details about the data which may (not) be processed in such environment
- the documentation about the cloud service, including the security recommendations from the CSP (and internal security requirements if available)
- other security requirements documents from PSA portal, because for applications hosted in the cloud, they still apply

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Disruption of availability
- Denial of executed activities
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-1/5.0

Req 2	The CSPs and the specific service models or services must not be used unless approved by DTAG Board of Management. The restrictions given in the approval related to approved data classes and information protection classes must be followed. In addition to BoM, other official procedures and processes, such as Simplified Cloud Approval Process (SCAP) for the cloud service approval may be used.
-------	---

The CSP is responsible for security of their infrastructure, services and APIs, which is basis for overall security. CSP is also responsible for fulfillment of their legal obligations and SLA. Since it would be impossible for all CSP customers to perform audits of a CSP, this is typically done by independent auditors which carry out audits and create audit reports.

The most relevant report from audit is SOC-2 type 2 report which needs to be evaluated by Telekom Security before signing the contract with CSP. If BSI C5 audit report is available, it must be evaluated too.

In addition to technical security evaluation of CSP based on audit reports, legal and contractual aspects are combined and recommendation is presented to DTAG Board of Management which in the end decides whether the usage of the particular CSP is allowed, along with the respective list of specific service models and specific services.

One may look at "Trusted Cloud Provider Approach" in YAM because the information there can be refreshed more often than in this document: <https://yam-united.telekom.com/pages/cloud-infrastructure-security/apps/blog/blog/view/d8f8b89f-9a94-4c40-a5aa-16fcbf6f1504>. Further information may be found also here at DT Cloud Taskforce: <https://yam-united.telekom.com/pages/dt-cloud-taskforce/apps/wiki/wiki/list>, where also contacts to responsible persons can be found. Link to fresh version of Simplified Cloud Approval Process (SCAP), especially useful for SaaS cloud services, can be found here: <https://yam-united.telekom.com/pages/psa/apps/content/materials>. One may find appropriate contacts on the same links to obtain information about already approved cloud services.

Motivation: Every service model and particular service of every CSP may have approval for different data classes, e.g. the approval for processing and storing personal data in one CSP's IaaS does NOT automatically mean that processing and storing of personal data is allowed for another CSP's IaaS even if in both cases IaaS is approved as a service model.

Due to the rising number of public cloud providers and their services, certain PaaS and SaaS providers or services which were not included in initial DTAG BoM approval can be allowed by other management entities, but such management entities must include appropriate DTAG Security responsible persons.

Implementation example: Example of an approval: Based on evaluation, DTAG BoM has decided that cloud provider X is allowed to be used by DTAG for all kinds of IaaS workloads with any kind of data, along with key-value database from the PaaS portfolio for all data except personal data, while all SaaS services are prohibited.

Can DTAG use VMs, virtual routers, firewalls and other services tightly related to VM compute services?

- Yes

Can DTAG store data inside key-value database?

- In general, yes, but details of an approval (fine print) states which kind of data.

Can DTAG store data inside relational database provided as PaaS from this CSP?

- No. Not in the list of approved services.

Can DTAG use IoT SaaS solution from CSP X for management of IoT devices?

- No. SaaS is not approved from CSP X.

Can DTAG use CSP X to work with or store PCI-DSS data?

- Not in general. Although many cloud providers are PCI-DSS compliant, it is necessary to check whether the certification applies to all services of CSP which DTAG would like to use. For use of PCI-DSS there is an additional requirement in this chapter.

Can DTAG process and store personal data in IaaS of CSP X?

- Yes.

Can DTAG store personal data in key-value database from CSP X?

- No.

Can DTAG use software from vendor Z which hosts their software on top of Giant Swarm on Azure?

- Vendor Z is considered as "3rd level" vendor which utilizes Giant Swarm, which is SaaS provided by Microsoft Azure. If Azure is approved as CSP and there is an approval for SaaS services including Giant Swarm on top of Azure, the answer would be generally yes. However, it is still necessary to perform security evaluation of the software from vendor Z, as much as possible for SaaS solutions, paying special attention to data classification restrictions. If there would be no approval to use Giant Swarm on Azure, a separate DTAG BoM decision (or equivalent process if existing) would be needed for usage of software from vendor Z.

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Unauthorized use of services or resources

ID: 3.66-2/5.0

Req 3	All applicable laws, regulations and corporate policies must be considered and followed prior to using public cloud for a system or data.
-------	---

No matter whether migrating existing system (or part of) or data to the public cloud, creating a new system which will work with new or existing data, it is necessary to verify that all laws, regulations and corporate policies are fulfilled. Some examples include:

- lawful interception
- data retention
- restrictions on geographical borders, both for data at rest (e.g. which regions, data centers or availability zones may be used), in use and in transit (e.g. is it allowed to transfer data over Internet or over geographical borders)
- telecommunications/communications act
- data provisioning

It is necessary to verify that the cloud provider offers sufficient controls to fulfill the requirements and that these requirements are implemented as policies for the system or data which is hosted in the public cloud.

Motivation: Not following laws and regulations may incur financial or reputation damage to DTAG and have legal consequences for employees.

Implementation example: In terms of implementation, cloud providers offer a variety of configurable policies which allow the cloud consumer to configure restrictions required by law, regulatory requirements and corporate policies throughout the whole landing zone, i.e. for all consumers of the cloud for a specific organization. For all restrictions which cannot be enforced through CSP's policies, appropriate administrative and/or organizational controls must be in place. People who use the cloud services need to be informed about the restrictions, learn them and follow them.

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Unauthorized use of services or resources

ID: 3.66-3/5.0

Req 4	The CSP and its services must not be used for data with need of protection, especially PCI-DSS data, unless the CSP provided appropriate guarantees like PCI-DSS certification and this type of usage was approved by DTAG Board of Management.
-------	---

Although many cloud providers are PCI-DSS compliant, it is necessary to check whether the certification applies to all services of CSP which DTAG would like to use. In addition to the certification, it is necessary that the DTAG Board of Management approved the CSP and specific services to be used with PCI-DSS data.

Motivation: Not following PCI-DSS requirements for handling credit card data may result in fines for DTAG.

Implementation example: This procedure (verification if the specific service of the cloud service provider) applies generally to data with need of protection, not only to the PCI-DSS data, although PCI-DSS is a very prominent example of such use cases.

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources

ID: 3.66-4/5.0

Req 5	Clear exit strategy must exist before onboarding a project or a system on a public cloud.
-------	---

In the traditional world, this problem was solved by dual (or multi) vendor strategy for avoiding vendor lock-ins, but the transition from vendor to vendor was possible because equipment was bought and owned by DTAG. This is not the case when using public cloud providers. Having 2 or more CSPs at disposal is not sufficient for dual (or multi) vendor strategy, it is also necessary to move impacted systems and data as well, either from the CSP to another CSP or from the CSP to own premises.

Contracts and court proceedings cannot be considered as a viable replacement for an exit strategy.

Examples of problems:

- CSP X raised prices which makes usage too costly (production cost of a service is higher than revenue).
 - CSP Y did not, so migration of services and data makes sense, but there must be a clear concept from the very beginning to avoid vendor lock-in.
- CSP went bankrupt, equipment is seized, and operations are terminated.
- CSP is providing bad service (not fulfilling SLA) and cannot solve technical issues for hours/days/months.
- CSP will discontinue service which DTAG is using within weeks/months (periods are typically much lower compared to EOL announcements for traditional software and hardware).
- CSP cannot be longer used due to the new laws and regulations.
- Developing application only around specific APIs from one CSP (e.g. AWS BeanStalk) would mean that the whole application needs to be rewritten for a migration to another CSP.

Motivation: Despite good intentions and contracts, it is possible that the CSP will not be able to fulfill the needs of DTAG and respective systems in the future. In order to account for potential legal, financial or technical problems which may arise in the future, it is necessary to have a clear plan how to get services and the data out of a specific cloud provider and at any given time.

Implementation example: If possible, cloud consumer should consider using more generic cloud services, so that transition to the public cloud from another CSP (or into own premises) is done with minimal modifications (or no modifications in best case scenario).

For this requirement the following threats are relevant:

- Disruption of availability
- Denial of executed activities

ID: 3.66-5/5.0

Req 6	The cloud services of the CSP must not be used unless done through the approved landing zone.
-------	---

CSP and DTAG typically sign a single contract and DTAG gets a single big account inside CSP. Inside that account, the required hierarchy can be created, e.g. mimicking organization structure. Basic rules (e.g. minimal password length policy) would apply to all DTAG users of this CSP, while some system/application specific (whether a VM can or cannot access the Internet) would be driven by owners of specific organization or a project.

Each exception to this rule requires a separate BoM (Board of Management) decision, according to the current situation at the time of writing.

Landing zone for SaaS cloud services is sometimes referred to simply as “administration portal” by SaaS cloud providers. This must be not confused with an existing DTAG landing zones for other clouds, such as DTIT landing zone for AWS, Azure and GCP.

Such landing zone of SaaS is typically smaller in terms of features compared to broader cloud offerings of hyperscalers, but still offers aspects like administrating users (adding and removing), permissions (setting up administrators, assignment of users to groups), restrictions about data retention (chat history, maximum lifetime of documents), log management (viewing and exporting audit logs), management of billing and payment, setting up quotas (limiting number of users or amount of data), viewing usage information, and setting up notifications for certain events.

Approval of SaaS cloud service by the security advisor also includes the approval of landing zone of the SaaS cloud service, since it is effectively the place for administration of this SaaS cloud service.

Motivation: Bypassing landing zone (e.g. by creation of another root account) would allow bypassing all the policies and is forbidden.

Implementation example: Use of only approved landing zones, i.e. do not separately create another landing zone without BoM approval.

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Unauthorized use of services or resources

ID: 3.66-6/5.0

Req 7	Unnecessary services must be disabled.
-------	--

After the installation of systems and software products, supplier-preset, local or network-accessible services are often active that are not required for the operation and functionality of the specific system in the intended operating environment.

However, in principle only the services actually required may be active on a system.

Accordingly, all services that are not required on a system must be completely disabled immediately after installation. It must be ensured that these services remain disabled even after the system is restarted.

In the context of the cloud this means: All cloud services which are not needed or out of scope for a landing zone or a project needs to be disabled as far as the cloud service provider supports this, e.g. by using service control policies.

In context of SaaS cloud services, this effectively translates to disable unwanted and unnecessary features and (sub)services.

Motivation: Active services that are not required unnecessarily increase the attack surface of a system and, as a direct consequence, the risk of a successful compromise. This risk can be further increased if - as is often observed with services that are not required - a targeted examination and optimization of the configuration with regard to security does not take place sufficiently.

Implementation example: For Azure or AWS you can use Azure policies or service control policies to deny access to the services which are not needed.

Examples for SaaS services include disabling file sharing if not in use, protocols and interfaces which are not needed, or analytics for translation services, especially if the data would be used to improve the translation service to prevent data leaks.

ID: 3.66-7/5.0

4. Governance and Management

Req 8	Process for handling abuse and incident notifications from and to the cloud service provider must be established.
-------	---

Cloud service providers typically notify customers on abuse and incidents which they detect, and sometimes even on identified irregularities (e.g. huge jump in resource consumption). The landing zone operators need to, in cooperation with responsible SOC, establish a process which enables handling of such notifications 24/7.

In the same way, it needs to be possible for a cloud user or landing zone operator to report incidents and abuse to the cloud service provider, which can be organized through the landing zone operators in agreement with responsible SOC.

Motivation: CSPs may detect certain abuse and incidents, even if the customer does not notice that, and will in such cases notify their customers.

Implementation example: Mailbox which is monitored 24/7 or phone number of 24/7 on-duty personnel.

For this requirement the following threats are relevant:

- Unauthorized use of services or resources
- Disruption of availability
- Denial of executed activities
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-8/5.0

Req 9	Central asset inventory must be maintained for all assets in the cloud and the process for management of assets must be established.
-------	--

Each system in the cloud is responsible to keep track of their assets. Assets should be tagged if possible, or marked in any other way so that it is possible to easily identify system owner and people who need to be notified in case of security incidents, along with the criticality, priority in case of recovery, and business relevance of the specific asset. Details will depend on the landing zone and the way how the security incident handling has been organized between landing zone operators, system owners and responsible SOC.

Motivation: All cloud assets (virtual machines, IP addresses, storage buckets, URLs of serverless functions, ...) typically come with a price tag (monthly), therefore it is extremely important to keep track of all the assets. Additional cost is not the only concern in public cloud, leftover assets can be also used as easier point for attacks.

Implementation example: If available, central asset management system should also include the asset information from the public cloud, preferably through the landing zone.

The best way a responsible SOC to access asset inventory in the cloud is through API because the information is then up-to-date, however, in some cases regular periodic exports and snapshots may serve the purpose if API access is not possible or convenient.

SaaS cloud services rarely offer detailed report of assets, but in some cases a billing statement or a similar periodically produced summary may be sufficient. On the other hand, more effort could be expected for populating asset inventory for a backup-as-a-service.

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Unauthorized use of services or resources
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-9/5.0

Req 10	Security tasks must be automated to reduce human errors during implementation.
--------	--

Typical tasks which must be automated whenever feasible are listed in Example section.

Motivation: In modern highly/fully automated environments like cloud (and CI/CD chains used by DevOps/DevSecOps organizations), manual processes only hinder fast pace which is expected and lead to even more common human/operator errors.

Implementation example: It is necessary to automate crucial processes:

- IAM related tasks - account, user, group and policy management
- security hardening of systems (security as code)
- security scanning and testing in CI/CD chains
- compliance and security checks of cloud services and systems running in the cloud (including verification of software/packages)
- vulnerability detection (both in CI/CD chains and in live environments)
- log data correlation and attack detection
- key, certificate and secret management
- detection and alerting of misconfigurations which cannot be enforced by policies
- security documentation generation for PSA process

SaaS cloud services typically require less configuration, i.e. have less options to configure. Although some of SaaS cloud services have configuration available only through web interface where no automation can be implemented at all, many of SaaS cloud services providers offer CLI and/or API access for configuration of some options, such as managing users, roles and permissions, exporting logs, managing quota or generating usage reports etc. Such CLI and API access can be used for automation of the security tasks.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Disruption of availability
- Denial of executed activities
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-10/5.0

Req 11	Features that are not required in the software, including application programming interfaces (APIs) and command line interface (CLI), and hardware must be deactivated.
--------	---

During the initial installation of software, features may have been activated by default that are not necessary for the operation and functionality of the specific system. Features are usually an integral part of the software that cannot be deleted or uninstalled individually.

Such features must be disabled immediately after the initial installation through the software's configuration settings, so that they remain permanently disabled even after the system is rebooted.

Even before delivery or during initial commissioning, features may have been activated by default in the hardware that are not required for the purpose of the specific system. Such functions, for example unnecessary interfaces, must also be permanently deactivated immediately after initial commissioning.

Motivation: A system's hardware or software often contains enabled features that are not being used. Such features can be an unnecessary target for manipulation. Furthermore, there is a potential that unauthorized access to areas or

data of the system can be created.

Implementation example: [Example 1]

Deactivation of debugging functions in the software that are used in the event of fault analysis, but do not have to be active during normal operation.

[Example 2]

Disabling unused network interfaces of a server.

ID: 3.66-13/5.0

Req 12	Software and hardware of the system must be covered by security vulnerability support from the supplier.
--------	--

Only software and hardware products for which there is security vulnerability support by the supplier may be used in a system.

Such support must include that the supplier

- continuously monitors and analyzes the product for whether it has been affected by security vulnerabilities,
- informs immediately about the type, severity and exploitability of vulnerabilities discovered in the product
- timely provides product updates or effective workarounds to remedy the vulnerabilities.

The security vulnerability support must be in place for the entire period in which the affected product remains in use.

Support phases with limited scope of services

Many suppliers optionally offer time-extended support for their products, which goes beyond the support phase intended for the general market, but is often associated with limitations. Some suppliers define their support fundamentally in increments, which may include limitations even during the final phase before the absolute end date of regular support.

If a product is used within support phases that are subject to limitations, it must be explicitly ensured that these restrictions do not affect the availability of security vulnerability support.

Open Source Software and Hardware

Open Source products are often developed by free organizations or communities; accordingly, contractually agreed security vulnerability support may not be available. In principle, it must also be ensured here that the organization/community (or a third party officially commissioned by them) operates a comprehensive security vulnerability management for the affected product, which meets the above-mentioned criteria and is considered to be reliably established.

Motivation: Hardware and software products for which there is no comprehensive security vulnerability support from the supplier pose a risk. This means that a product is not adequately checked to determine whether it is affected by further developed forms of attack or newly discovered vulnerabilities in technical implementations. Likewise, if there are existing security vulnerabilities in a product, no improvements (e.g. updates, patches) are provided. This results in a system whose weak points cannot be remedied, so that they remain exploitable by an attacker in order to compromise the system or to adversely affect it.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Disruption of availability
- Denial of executed activities
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.01-1/9.0

Req 13	Known vulnerabilities in the software or hardware of the system must be fixed or protected against misuse.
--------	--

Known vulnerabilities in software and hardware components must be fixed by installing available system updates from the supplier (e.g. patches, updates/upgrades). Alternatively, the use of workarounds (acute solutions that do not fix the vulnerability, but effectively prevent exploitation) is permissible. Workarounds should only be used temporarily and should be replaced by a regular system update as soon as possible in order to completely close the vulnerabilities.

Components that contain known, unrecoverable vulnerabilities must not be used in a system.

The treatment of newly discovered vulnerabilities must also be continuously ensured for the entire deployment phase of the system and implemented in the continuous operating processes of security patch management.

Motivation: The use of components without fixing contained vulnerabilities significantly increases the risk of a successful compromise. The attacker is additionally favored by the fact that, as a rule, not only detailed information on vulnerabilities that have already become known is openly available, but often also already adapted attack tools that facilitate active exploitation.

Implementation example: Following the initial installation of an operating system from an official installation medium, all currently available patches and security updates are installed.

Additional information:

The primary sources of known vulnerabilities in software/hardware are lists in the release notes as well as the security advisories from the official reporting channels of the supplier or independent CERTs. In particular, the reporting channels are sensibly integrated into continuous processes of security patch management for a system, so that newly discovered vulnerabilities can be registered promptly and led into operational remedial measures.

As a complementary measure to the detection of potentially still contained types of vulnerabilities that have in principle already become known, targeted vulnerability investigations of the system can be carried out. Particularly specialized tools such as automated vulnerability scanners are suitable for this purpose. Examples include: Tenable Nessus, Qualys Scanner Appliance.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Disruption of availability
- Denial of executed activities
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.01-10/9.0

Req 14	A configuration repository must be used as a single source of truth for all configuration files and infrastructure-as-code (IaC) scripts that define the cloud resources, and such configuration repository must not be hosted inside the same cloud.
--------	---

The establishment of a configuration repository as the single source of truth for all configuration files and infrastructure-as-code (IaC) scripts is essential to ensure efficient, organized, and reliable management of cloud resources. By having a dedicated repository with version control, teams can easily track changes, maintain consistency, enhance collaboration, and support disaster recovery, contributing to a streamlined and transparent cloud infrastructure management process.

The repository must be protected with appropriate access controls to enforce the "need-to-know" and least privilege principle, allowing only authorized personnel access to specific configurations and the deployment to the cloud.

While some of SaaS cloud service providers have configuration available only through web interface where no automation can be implemented at all, many of SaaS cloud service providers offer CLI and/or API access for configuration. Such CLI and API access can be used for the configuration.

The configuration repository used for controlling the configuration of the cloud must be placed outside of the cloud which is being controlled so that it is possible to restore the cloud environment by using such configuration repository.

Motivation: Having a dedicated config repository, helps reduce configuration errors, streamline configuration management processes, enhance security by enforcing standardized configurations, and simplify troubleshooting and auditing activities.

Hosting a configuration repository of the cloud environment inside the same cloud environment could create a situation where a cloud environment cannot be repaired because the configuration repository is unavailable or unreachable, and configuration repository can be made available only if the cloud environment is repaired.

Implementation example: Infrastructure as code repository of AWS landing zone can be stored in neither AWS landing zone nor a CI/CD system hosted inside the same AWS landing zone. Such repository should be rather hosted on DTAG premises. Or in some special cases, such repository needs to have a fully functional mirror outside of the environment which it defines.

For this requirement the following threats are relevant:

- Disruption of availability
- Denial of executed activities
- Unnoticeable feasible attacks

ID: 3.66-20/5.0

Req 15	Tags with appropriate and meaningful metadata must be applied to all cloud resources (virtual machines, storage buckets, databases, network components, etc.) that support tags.
--------	--

Tagging is an essential feature that must be implemented in cloud environments to organize resources within an organization and reduce complexity and administrative overhead. Enforcing a tagging policy can ensure consistency, accuracy, and standardization of all resources in a cloud environment, allowing for better resource tracking and management.

Some SaaS cloud services offer tagging, especially those which handle various kinds of data like cloud storage or cloud backup. Tags can be used here to represent data classification and ownership.

Motivation: Tags enable the categorization of resources with own metadata. This can be useful to implement e.g. security configurations in an automated manner and to classify data.

Implementation example: The specifications of which tags to use are defined by the landing zone. Additional tags can be defined by applications themselves.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-21/5.0

Req 16	Permissions and controls implemented as the landing zone must be verified before use and after significant changes.
--------	---

Verification of the correctness of controls implemented as the landing zone is mandatory before use and whenever there are substantial changes in the landing zone (e.g. complete redefinition of roles). This applies to at least the main landing zone and to the parts of the landing zone which have diverging (overriding) policies compared to the top-level landing zone. The security advisor must be contacted and the application of the changes must be aligned with the se-

curity advisor.

The usage of automated tools for penetration tests and CIS benchmarks can become very expensive in no time, so it is very important to appropriately set the right scope of tests and scans according to the needs.

Motivation: Any change and of configuration and controls (especially policies in terms of cloud landing zone) may result in new security vulnerabilities.

Implementation example: The following tools and methods can be used for performing the verification:

- automated penetration test tools (most common example is privilege escalations)
- custom (manual) penetration tests - typically require a creative approach
- CIS benchmarks (<https://www.cisecurity.org/>) and similar automated verification/reporting tools

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Disruption of availability
- Denial of executed activities
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-22/5.0

4.1. Session Protection

Session protection is essential to prevent unauthorized access, session hijacking, and maintain data confidentiality for user interactions with webservices. In the cloud environment especially the webportal of the cloud is the place to check for session protection.

Req 17 Sessions must be protected against unauthorized takeover ("session hijacking").

Interfaces that provide session functionality to the system must implement technical measures to prevent a legitimate user's session from being taken over and continued by an unauthorized third party.

Such protection can be achieved, for example, by implementing a combination of the following options that makes sense for the specific system:

- At the transport layer: Use of the TCP protocol (with its sequence numbers) and corresponding filter lists
- At the session layer: Use of the TLS Protocol
- At the application layer: Negotiation of a random secret session key between sender and receiver to authorize all session traffic (e.g. session ID, session cookie, session token)
- Use of cryptographic methods to protect session keys from eavesdropping or modification attacks

Motivation: Unprotected sessions can potentially be hijacked and continued by an attacker in order to exercise unauthorized access to the system in the context of the affected user.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.01-16/9.0

Req 18 The system must allow users to log out of their current session.

The system must have a feature that enables the logged-in user to log out at any time. It must not be possible to resume a logged-out session without re-authenticating the user.

Motivation: A user must retain complete control over the sessions he has established in order to be able to terminate his access to a system at any time according to the situation and thus protect data and functions exposed via this access. In addition, the user must be able to assume that sessions specifically terminated by him cannot subsequently be resumed and continued by unauthorized third parties.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.01-17/9.0

Req 19 Sessions must be automatically terminated after a period of inactivity adapted to the intended use.

It is necessary that sessions on a system are automatically terminated after a specified period of inactivity.

For this reason, a time-out for sessions must be set. The time period to be selected here depends on the use of the system and, if applicable, the physical environment. For example, the time-out for an application in an unsecured environment must be shorter (a few minutes) than the time-out for an application used by operations personnel for system monitoring tasks in an access-protected area (60 minutes or more).

Motivation: For an open but unused session, there is a risk that an illegitimate user may take over and continue it unnoticed in order to exercise unauthorized access to the system and the data contained therein on behalf of the affected user.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.01-18/9.0

5. Authentication and Authorization

In order to make use of approved CSP through a proper contract, it is necessary to create "Landing Zone" for every CSP. Landing zone is a concept which allows central management of identities and projects inside CSP for the whole DTAG group.

Landing zone starts with "root account" for whole DTAG, managed by team responsible for managing CSP usage (cloud management team in rest of the chapter). The term "root account" is also known simply as administrator in certain SaaS cloud services. The cloud management team manages set of basic (security) policies, users and projects. Each project/system/department/organization would get access to CSP over the landing zone and they would be able to consume cloud services with predefined quota limits and within basic security policies. Each organization has a freedom to customize their portion under given initial constraints. Each project under every organization would be able to customize further.

Landing zone is prerequisite for complete and comprehensive security monitoring and accounting.

Req 20	The use of system functions that require protection as well as access to internal or confidential data must not be possible without prior authentication and authorization.
--------	---

The use of functions of the system that require protection as well as access to data classified as internal or confidential must only be possible after the user has been uniquely identified and successfully authenticated by means of a user account. A user account consists of a unique user name and associated authentication features. In addition, it must be verified that the user is authorized to access the affected functions and data within the user role assigned to him or her in the system.

An exception to this are functions and data that may be used publicly without restriction; for example, the area of a website on the Internet where only public information is provided.

Examples of features that require prior authentication include:

- Remote access to network services (such as SSH, SFTP, web services)
- Local access to the management console
- Local use of operating system and applications

This requirement also applies without restriction to any machine access to the system (Here, implementation is usually carried out through the use of so-called "technical user accounts", also known as "machine-2-machine" or "M2M" accounts).

Motivation: The unambiguous authentication and authorization of access to a system are elementary to protect functions and data from misuse.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.01-19/9.0

Req 21	User accounts must be protected with at least two authentication features from different factors.
--------	---

All user accounts of the system must be protected from unauthorized use.

This is to be achieved by using authentication features that allow the legitimate user to be uniquely identified.

Authentication features are differentiated into superordinate groups, so-called "factors":

- Factor KNOWLEDGE: "something that only the legitimate user knows"

- Factor POSSESSION: "something that only the legitimate user possesses"
- Factor INHERENCE: "something that only the legitimate user is"

At least two authentication features must be combined.

These must come from different factors ("Multi-Factor Authentication"; short: "MFA").

A combination of authentication features of the same factor (example: two different PINs) is not allowed.

Examples of acceptable authentication features

- from the KNOWLEDGE factor: PINs, passphrases (passwords*)
- from the POSSESSION factor: smart cards, FIDO 2 Level 2 security keys, device-bound passkeys, cryptographic keys stored in a TPM/HSM, OTPs (One-Time Passwords) generated in an authenticator app that has been approved for use in Deutsche Telekom
- from the INHERENCE factor: biometric features such as fingerprints, iris or retina scans, face recognition

*) Important note on the use of passwords

Passwords must be strictly avoided as an authentication feature.

They are only acceptable if more secure alternatives cannot be implemented.

Authentication features that can be easily forged or manipulated are generally not permitted.

These include:

- Phone numbers and caller IDs
- IP addresses
- Geographical locations
- VPN affiliations and network information
- Status of clients, such as workstation systems or mobile devices

Note on technical user accounts

Technical user accounts used for machine access (also known as "machine-to-machine" or "M2M" accounts) are exempt from the obligation to use multi-factor authentication, as this is usually not feasible there.

Persons are not supposed to use non-interactive access directly (like API access), but rather to use technical user (sub)accounts which belong to them. Technical user accounts are covered later in the document.

Motivation: User accounts that are not protected by appropriate authentication features can be abused by an attacker to gain unauthorized access to a system and the data and applications stored on it. Due to a variety of common attack methods, the protection of user accounts by a single authentication feature is now considered weak, especially if the classic password is used here. A combination of authentication features from different factors significantly increases security for user accounts.

Implementation example: Examples of MFA

MFA Combination	Description	Application Note
Telekom "myCard" plus PIN	1st Factor: Certificate implemented in hardware ("Possession") 2nd Factor: PIN ("Knowledge")	(none)

FIDO2 Level 2 Security Key plus fingerprint	1st Factor: Crypto key stored in hardware ("Possession") 2nd Factor: Biometric feature ("Inherence")	Biometric feature must be pre-configured, otherwise the biometric sensor will only work as a "button" and does not provide the 2nd factor ("Inherence").
Device-bound Passkey in WH4B* on Nucleus PC plus PIN	1st Factor: Crypto key stored in hardware ("Possession") 2nd Factor: PIN ("Knowledge")	(none)
Device-bound Passkey in Secure Enclave on iOS Device plus FaceID	1st Factor: Crypto key stored in hardware ("Possession") 2nd Factor: Biometric feature ("Inherence")	(none)
Password plus OTP App approved for use (e.g. MS Authenticator)	1st Factor: Password ("Knowledge") 2nd Factor: specific generated One Time Password ("Possession")	Passwords are only eligible if more secure alternatives cannot be implemented. The synchronization of the authenticator app via a cloud service as well as sharing (e.g. via a password manager) the OTP is not permitted.

*) WH4B: "Windows Hello for Business"

ID: 3.66-28/5.0

Req 22 Privileged user accounts must be protected with at least two authentication features that come from different factors and provide increased resilience against phishing attacks.

A privileged user account is a user account that has advanced permissions within a system. Advanced permissions allow access to configuration settings, features, or data assets that are not available to regular users of the system. In direct dependence on the specific tasks performed by means of a privileged user account within a system, the advanced permissions assigned may be specifically limited or may include fully unrestricted system access.

Examples of privileged user accounts:

- Accounts for administration, maintenance, or troubleshooting tasks
- Accounts for user administration tasks (e.g. creating/deleting users; Assign permissions; password resets)
- Accounts that are authorized to legitimize, trigger, or stop business-critical operations
- Accounts with access to data classified as SCD (Sensitive Customer Data) in the interest of the Group Deutsche Telekom, its customers or the public
- Accounts that have extensive access to data defined as "personal" according to the EU-GDPR (e.g. mass retrieval of larger parts or the entire database)

To protect privileged user accounts, at least two authentication features must be combined. These must come from different factors ("KNOWLEDGE", "POSSESSION" or "INHERENCE"). A combination of authentication features of the same factor is not permitted. At least one of the authentication features must not be vulnerable to phishing attacks.

In a phishing attack, the legitimate user is trickily manipulated to reveal his login data to the attacker, for example by means of fake login pages. A phishing attack can also be carried out against a combination of multiple authentication

features.

Phishing-resilient authentication features prevent attacks based on their technical implementation. These authentication features use cryptographic methods and create a strong bond between the user, the device and the service. They also require manual user interaction during authentication, which cannot be automated; for example, to submit a fingerprint scan to a hardware token.

Considered not vulnerable to phishing attacks are

- FIDO2 Level 2 Security Keys
- Certificates implemented in hardware

Important note on the use of passwords

Passwords must be strictly avoided as an authentication feature.

They are only acceptable if more secure alternatives cannot be implemented.

Note on technical user accounts

Technical user accounts used for machine access (also known as "machine-to-machine" or "M2M" accounts) are exempt from the obligation to use multi-factor authentication, as this is usually not feasible there.

Motivation: Privileged user accounts pose an increased risk to the security of a system. If an attacker successfully compromises such a user account, they will be granted broad privileges that allow them to take control of the system or system parts, disrupt system functions, view/manipulate processed data, or influence business-critical operations. The combination of multiple authentication features of different types already significantly minimizes the risk of a user account being compromised, but for privileged user accounts, special attention must be paid to resilience against phishing attacks. The phishing attacks that can be observed are becoming increasingly tricky and difficult for users to detect, so that multifactor authentication without increased resilience against phishing can no longer be considered sufficient for privileged user accounts.

Implementation example: Examples of MFA with increased resilience to phishing attacks

MFA Combination	Description	Application Note
Telekom "myCard" plus PIN	1st Factor: Certificate implemented in hardware ("Possession") 2nd Factor: PIN ("Knowledge")	(none)
FIDO2 Level 2 Security Key plus fingerprint	1st Factor: Crypto key stored in hardware ("Possession") 2nd Factor: Biometric feature ("Inherence")	Biometric feature must be pre-configured, otherwise the biometric sensor will only work as a "button" and does not provide the 2nd factor ("Inherence").
FIDO2 Level 2 Security Key plus PIN	1st Factor: Crypto key stored in hardware ("Possession") 2nd Factor: PIN ("Knowledge")	(none)

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.01-21/9.0

Req 23	Default (root) account(s) must be treated as privileged user accounts, and where possible, following four-eyes principle.
--------	---

Root account for the whole DTAG inside CSP is created with all available privileges. Such account is a security risk and must not be used for daily operations. Other accounts with extremely restricted set of privileges need to be used for daily operations. However, in some rare cases, it is necessary to perform tasks with the root account, and therefore this type of an account must be treated as a privileged user account and treated accordingly (like MFA, increased resilience against phishing attacks etc.).

It is a very good practice to monitor if e-mail address (account identifier in general) is included in case of breaches at the CSP.

All above mentioned credentials and devices for such root accounts must be securely stored in a secure location, like physical safe. It is recommended to store them on 2 locations (e.g. password + token at one location, recovery secrets on another location) for redundancy because earthquakes, fires and floods still happen.

Every action of default/root account(s) ("superadmins") must be fully logged for auditing purposes. It is advisable to also configure informing of all individuals involved in the top-level cloud administration (e.g. via e-mail) whenever a login to such account occurs, or in case of some other changes to such account(s).

This requirement is in line with Requirement 16 from 3.63. ("Operations in DT IT") for those systems where 3.63. applies, however, it is necessary to note that password length/complexity from this requirement overrides Requirement 16 from 3.63.

Motivation: Reduction of the attack surface.

For this requirement the following threats are relevant:

- Unauthorized access to the system

ID: 3.66-30/5.0

Req 24	Technical user accounts must use appropriate authentication mechanisms (short lived tokens or access keys instead of MFA).
--------	--

Technical user accounts are necessary for automation and they cannot use typical MFA which is mandatory for persons. Fortunately, there are other options which can be used for automated M2M communication scenario. One such options is use of tokens, access keys and certificates (later, only the term token is used) which have very limited scope, privileges and lifetime. Depending on the selected CSP, options might vary but the following restrictions need to be applied whenever possible:

- Technical user accounts or their access keys/tokens are not allowed to be used by persons
- Tokens for technical user accounts must have limited lifetime and periodically exchanged.
 - If there is no way to specify initial lifetime of a token, there must be an established process for rotation of older tokens.
 - Process for rotation of tokens needs to consider smooth transition from old to new token.
 - Token lifetime needs to be appropriately monitored to avoid downtime caused by expired tokens.
 - Stolen or leaked tokens must be immediately revoked/deleted.

Motivation: Short lived credentials help to minimize the impact of potential security breaches, limit the exposure of information with need of protection, and maintain a proactive security posture.

Implementation example: Use OpenID Connect with GitLab to obtain temporary technical user account credentials for the CSP.

For this requirement the following threats are relevant:

- Unauthorized access to the system

- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.66-31/5.0

Req 25	Secret (key, token, certificate) rotation must be frequent. State-of-the-art secret protection mechanisms must be used.
--------	---

Secrets used in a public cloud requirement are at higher risk of being stolen or leaked. To minimize the impact of stolen secrets, they need to be generated with shorter lifetime compared to traditional deployments. It is practically impossible to predefine specific lifetime of secrets for every possible system and data class, therefore the final decision is left to agreement between application owner and responsible security officer. The currently valid security requirements for secrets are contained in PSA document 3.50 "Cryptographic Algorithms and Security Protocols". As of this writing, a rotation of secrets is necessary every 24 months or on indication of compromise. Rotation of secrets needs to be implemented in a way which allows smooth transition, i.e. without service interruption.

Secrets will be used in the public cloud environment itself, so it is necessary to employ state-of-the-art secret protection mechanisms. One such example is protection of private keys at rest in RAM against speculation and memory channel attacks like Spectre, Meltdown, Rowhammer and RAMBleed. More information about this specific protection mechanism can be found at: <https://marc.info/?l=openbsd-cvs&m=156109087822676&w=2>.

Examples:

- Example for HTTPS certificate duration installed on front-end web servers or proxies: validity of a certificate for more than 1 year might be too high. 3 months would be probably a better choice.
- Private-public key pair in asymmetric encryption (RSA key pairs for SSH) duration for more than 2 years might not be the best idea (and the rotation is mandatory), even with protection of private keys in RAM offered by SSH software.
- Some further typical examples can be found at <https://www.keylength.com/>.

This requirement is in line with Requirement 16 from 3.63. ("Operations in DT IT") for those systems where 3.63. applies.

Motivation: Regularly rotating secrets help minimize the impact of potential security breaches, limit the exposure of information with need of protection, and maintain a proactive security posture. Security breaches are, in most cases, exposed secrets through direct use in the source code that is visible to unauthorized persons through a cross-code repository (Gitlab, Github, etc.), or the propagation of the keys through other unintended actions.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-32/5.0

Req 26	Predefined authentication attributes must be changed.
--------	---

After the takeover or initial installation of a system, there are usually predefined authentication attributes (e.g. passwords, SSH keys, SSL/TLS Certificates) in the system, as assigned by manufacturers, developers, suppliers or automated installation routines.

Such predefined authentication attributes must be changed to new, individual values immediately after the takeover or installation of the system.

In the context of the cloud this means:

Also in the cloud, the creation or initiation of a new service can include a standard/predefined authentication attribute, for example a standard password for a service user. As mentioned above these attributes must be changed to a new individual value after provisioning the service in the cloud.

Administrators or users of cloud services, especially SaaS, may get initial password via email, and in such case, they are obliged to change the password after logging in for the first time.

Motivation: Values predefined by third parties in authentication attributes cannot be trusted because they do not represent a controlled secret. Affected authentication attributes can be misused by unauthorized persons to access and compromise systems. This risk is significantly increased if commonly known default values are used for authentication attributes (e.g. a default password for the administrator user account in a particular software product).

ID: 3.66-33/5.0

Req 27	User accounts must ensure the unique identification of the user.
--------	--

Users must be identified unambiguously by the system.

This can typically be reached by using a unique user account per user.

So-called group accounts, which are characterized by the fact that they are used jointly by several people, must not be used. This also applies without restriction to privileged user accounts. Most systems initially have only a single user account with administrative privileges after the basic installation. If the system is to be administered by several persons, each of these persons must use a personal, individual user account to which appropriate administrative authorizations or roles are assigned

A special feature are so named technical user accounts. These are used for the authentication and authorization of systems among themselves or of applications on a system and can therefore not be assigned to a specific person. Such user accounts must be assigned on a per system or per application basis. In this connection, it has to be ensured that these user accounts can't be misused.

Ways to prevent misuse of such user accounts by individuals include:

- Configuration of a password that meets the security requirements and is known to as few administrators as possible.
- Configuring the user account that only a local use is possible and a interactive login isn't possible.
- Use of a technique for authentication of the specific user account with public and private key or certificates.
- Limiting the access over the network to legitimate systems.

Additional solution must be checked on their usability per individual case.

Motivation: Unambiguous user identification is mandatory to assign a user permissions that are necessary to perform the required tasks on the system. This is the only way to adequately control access to system data and services and to prevent misuse. Furthermore, it makes it possible to log activities and actions on a system and to assign them to individual users.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.01-22/9.0

Req 28	The permissions for users and applications must be limited to the extent necessary to fulfill their tasks.
--------	--

The permissions on a system must be restricted to such an extent that a user can only access data and use functions that he needs in the context of his work. Appropriate permissions must also be assigned for access to files that are part of the operating system or applications or that are generated by the same (e.g. configuration and logging files).

In addition to access to data, applications and their components must also be executed with the lowest possible permissions. Applications should not be run with administrator or system privileges.

Motivation: If a user is granted too far-reaching permissions on a system, he can access data and applications to an extent that is not necessary for the fulfillment of the assigned tasks. This creates an unnecessarily increased risk in the event of abuse, in particular if the user or his user account is compromised by an attacker.

Applications with too far-reaching permissions can be misused by an attacker to gain or expand unauthorized access to data and system areas in need of protection.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.01-23/9.0

Req 29	Technical user accounts must have limited privileges. API access for technical user accounts must be limited to minimum which is necessary.
--------	---

Permissions for technical user accounts must be restricted to such an extent that technical user accounts only access data and use features that are needed in the course of their work.

That is very important, because technical user accounts cannot use typical MFA as an additional security mechanism, which would give them an additional protection.

Depending on the selected CSP, options might vary but the following restrictions need to be applied whenever possible:

- Technical user accounts must have no more than bare minimum of privileges needed for tasks they perform
 - Although this is a valid requirement for all user accounts, technical user accounts are typically used by automation tools which have predefined API calls and therefore do not need more privileges than required by possible API calls.
- Technical user accounts must have limited scope for resource access (only a subset of resources)
- Technical user accounts or their access keys/tokens are not allowed to be used by persons
- Tokens for technical user accounts must have limited lifetime and periodically exchanged.
- Access to CSP APIs must be limited in terms of network reachability.
 - APIs of CSP are available over Internet, however, it is possible in some cases to restrict network access for technical user accounts to APIs so that connections from only specific IP address ranges or locations is possible.

Motivation: If the rights granted to a user/service on a system are too broad, it could be possible to access data and applications for which viewing or the use is not permitted. This would give the opportunity to disclose or modify confidential data and to manipulate system files.

Implementation example:

- If technical user account needs only to read object storage, it must not have a privilege to modify, write or delete privileges for object storage.

- If technical user account is used for storing logs to object storage, it needs to have appropriate create/write privileges, but must not have privileges for reading, modifying or deletion of objects.
- If technical user account is used for provisioning of virtual machines, it must not have privileges for modification of security groups.
- If technical user account is used for management of only a subset of virtual machines (e.g. VMs with a specific tag), it must not have possibility to manage other virtual machines (VMs with different tags).
- If technical user account accesses CSP's APIs from internal networks (e.g. from VMs inside cloud) and not over Internet, and if CSP allows restriction of API access for given technical user account, Internet access for such technical user account must be removed.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources

ID: 3.66-36/5.0

Req 30	All public cloud user accounts must be under root account. Other root accounts must not be created.
--------	---

As specified previously, usage of CSP is allowed only through the proper landing zone, and since the landing zone is created by initial default root account, all subsequently created users must belong to that main account.

Creating another personal or business root account on that CSP would create a new landing zone for that user, which would allow bypassing security policies set for the previously described landing zone and therefore is forbidden.

Examples:

My project is special, and I need to use X and we must have a separate root account.

- A new BoM decision is necessary.

Motivation: Creating another personal or business root account on that CSP would create a new landing zone for that user, which would allow bypassing security policies set for the previously described landing zone and therefore is forbidden.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Disruption of availability
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-37/5.0

Req 31	IAM solution of DTAG must be used for user authentication if supported by CSP. If not, automation for managing identities in the cloud must be in place.
--------	--

Double administration of users brings issues like untimely creation and deletion of accounts (which can lead to data leaks), so it is necessary that accounts in the cloud are mapped to accounts in own DTAG's IAM. In terms of technologies - Entra ID (previously known as AAD or Azure Active Directory), SAML 2.0, OpenID+OAuth (2.0) and LDAPS (inside VPN connection only) are viable options.

In an unlikely case of not being able to integrate DTAG's own IAM solution for managing identities in the cloud, it is necessary to automate account management. Such automation for management of identities has to fulfill the "IAM (Identity Access Management) - Framework" (3.69.) requirements.

Preferred solution would be session based since no employee personal data would be exchanged with the cloud provider.

It is worth mentioning that this requirement also applies to applications running in the cloud, however, that is already covered with other security requirements (for operating systems and applications).

Example:

- Employee has left DTAG and his account in the cloud must be deleted to prevent him accessing the cloud resources.

In case of a system where 3.63. ("Operations in DT IT") applies, Requirement 7 from 3.63. applies as more specific requirement.

Motivation: The use of two or more IAM systems for one solution leads to a high level of complexity, increases the risk of inconsistencies and security gaps, compared to one IAM system. Always use already approved standard systems, to remove complexity and increase security.

Implementation example: One can find more information on the following links:

- Onboarding Guide telekom.de Entra ID – Tenant: <https://wiki.telekom.de/pages/viewpage.action?pagelId=1563181642>
- Information Collaboration Network: <https://yam-united.telekom.com/pages/a66ec908-117c-4fa8-94f3-502795da60c5/apps/wiki/abd69725-2c88-4187-8499-9ed6847cd0f8/list/view/5a395001-6a94-43ca-8770-68e920222a52>

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources

ID: 3.66-39/5.0

Req 32 Privileges for cloud usage must not be associated directly with users or services.

Terminology may vary between CSPs, so most common one was used here.

Privileges for cloud service consumption must be assigned to roles and never to individual users.

Users need to be organized in groups, and these group names must clearly identify their meaning (e.g. administrators, operators, reporters). Multiple users may share the same group.

Groups (or policies) allow mapping of between users and roles (set of privileges). In some instances, this allows users who have multiple roles to select an appropriate role (getting all necessary privileges) for the work.

This concept simplifies assignment and removal of privileges to specific users and allows better control over privileges, and even allows users switching from one area of responsibility to another one (e.g. developer -> operations).

Motivation: Assigning permissions to roles or groups offers benefits such as simplified management, consistency, granular control, enhanced security, scalability, and improved auditability, making it a best practice for managing access controls in cloud environments.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized use of services or resources

ID: 3.66-40/5.0

Req 33 Privileged Access Workstations (PAWs) must be used for administrative Tasks.

A Privileged Access Workstation (PAW) is a isolated operating system (hardware/virtualized) used for the purpose of securely accessing privileged accounts and resources.

A PAW must be used when performing administrative Task on a Public Cloud.

PAWs usually prohibit the use of risky applications like e.g. Web Browsers and Email Clients. They idea behind a PAW is to disconnect daily activities like e.g answering mails from administrative tasks.

Motivation: PAWs provide enhanced security for administrators working with high risk resources vulnerable to compromise, due to the use of trusted hardware and hardened dedicated systems.

Implementation example: Use a PAW solution for administrative tasks.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Denial of executed activities
- Unnoticeable feasible attacks

ID: 3.66-41/5.0

6. Authentication attribute "password"

Req 34	If a password is used as an authentication attribute, it must have at least 12 characters and contain three of the following categories: lower-case letters, upper-case letters, digits and special characters.
--------	---

A system may only accept passwords that comply with the following complexity rules:

- Minimum length of 12 characters.
- Comprising at least three of the following four character categories:
 - lower-case letters
 - upper-case letters
 - digits
 - special characters

The usable maximum length of passwords shall not be limited to less than 25 characters. This will provide more freedom to End Users when composing individual memorable passwords and helps to prevent undesired behavior in password handling.

When a password is assigned, the system must ensure that the password meets these policies. This must be preferably enforced by technical measures; if such cannot be implemented, organizational measures must be established. If a central system is used for user authentication [see also Root Security Requirements Document[i] "3.69 IAM (Identity Access Management) - Framework"], it is valid to forward or delegate this task to that central system.

Permissible deviation in the password minimum length

Under security-related criteria, conditions can potentially be identified for a system that enable the minimum password length to be reduced:

- It is generally permissible to reduce the minimum password length for systems that use additional independent authentication attributes within the authentication process in addition to the password (implementation of 2-Factor or Multi-Factor Authentication).
- Any reduction in the minimum password length must be assessed individually by a technical security advisor (e. g. a PSM from Telekom Security) and confirmed as permissible. In the assessment, the surrounding technical, organizational and legal framework parameters must be taken into account, as well as the system-specific protection requirements and the potential amount of damage in the event of security incidents.
- The absolute minimum value of 8 characters length for passwords must not be undercut.

Motivation: Passwords with the above complexity offer contemporary robustness against attacks coupled with acceptable user friendliness. Passwords with this level of complexity have proven their efficiency in practice. Trivial and short passwords are susceptible to brute force and dictionary attacks and are therefore easy for attackers to determine. Once a password has been ascertained it can be used by an attacker for unauthorized access to the system and the data on it.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.01-26/9.0

Req 35	If a password is used as an authentication attribute for technical accounts, it must have at least 30 characters and contain three of the following categories: lower-case letters, upper-case letters, digits and special characters.
--------	--

Technical user accounts are characterized by the fact that they are not used by people. Instead, they are used to authenticate and authorize systems to each other or applications on a system.

A system must only use passwords for technical user accounts that meet the following complexity:

- Minimum length of 30 characters
- Comprising at least three of the following four character categories:
 - lower-case letters
 - upper-case letters
 - digits
 - special characters

Motivation: Due to their use in machine-to-machine (M2M) communication scenarios, technical user accounts are often equipped with privileges that can be of high interest to an attacker to compromise infrastructures. Without mechanisms of extensive compromise detection, the risk of a password being determined or broken by an attacker can increase significantly over time. A significant increase in password length counteracts these risks and can also be implemented particularly easily in M2M scenarios, since handling a very long password is not a particular challenge for a machine (as opposed to a person).

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities

ID: 3.01-27/9.0

Req 36	If a password is used as an authentication attribute, the reuse of previous passwords must be prevented.
--------	--

A history of the previously used passwords must be recorded for each user account. When a password change is initiated for a user account, the new password must be compared with this password history. If the reuse of a password is detected, the password change must be rejected. This validation process must be implemented in the system on the basis of technical measures. If a central IAM system is used for user authentication, the implementation can be forwarded to the central IAM system or outsourced there [see also Root Security Requirements Document[i] "3.69 IAM (Identity Access Management) - Framework"].

In general, the password history should ensure that a password that has already been used can never be used again.

However, due to technical limitations, a password history cannot be recorded indefinitely in many IT/NT products. In this case, the following basic rules must be observed:

- a password that has already been used must not be reusable for a period of at least 60 days (measured from the point in time at which the affected password was replaced by another)
- in systems in which the period of at least 60 days cannot be implemented, the longest possible period must be configured. In addition, it must be confirmed by a Project Security Manager (PSM) that the configured period is still sufficient in the overall context of the system with regard to the security requirement.

Annotation:

Some IT/NT products do not offer any technical configuration parameters with which the password history can be linked directly to a time period, but only allow the definition of the number of passwords to be recorded. In such cases, the time period can alternatively be ensured by linking the following, usually generally available configuration parameters. Within the resulting policy, a user can only change his password once a day and, due to the number of passwords recorded, can reuse an old password effectively after 60 days at the earliest.

- Minimum Password Age: 1 day
- Password History: Record of the last 60 passwords used

With this implementation variant, it should be noted that the minimum age for the password should not be more than one day in order not to inappropriately restrict the user with regard to the fundamental need to be able to change the password independently at any time.

Motivation: Users prefer passwords that are easy to remember and often use them repeatedly over long periods of time when the system allows. From the user's point of view, the behavior is understandable, but effectively leads to a considerable reduction in the protective effect of this authentication parameter. With adequate knowledge of the user or information obtained from previous system compromises, an attacker can gain access to supposedly protected user accounts. Particularly in situations in which new initial passwords are assigned centrally as part of an acute risk treatment, but users change them immediately to a previous password for the sake of simplicity, there is a high risk that an attacker will resume illegal access. It is therefore important to prevent users from reusing old passwords.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.01-28/9.0

Req 37	If a password is used as an authentication attribute, users must be able to independently change the password anytime.
--------	--

The system must offer a function that enables a user to change his password at any time.

When an external centralized system for user authentication is used, it is valid to redirect or implement this function on this system.

Motivation: The fact that a user can change his authentication attribute himself at any time enables him to change it promptly if he suspects that it could have been accessed by a third party.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Denial of executed activities
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.01-29/9.0

Req 38	If a password is used as an authentication attribute, it must be changed after 12 months at the latest.
--------	---

The maximum permitted usage period for passwords is 12 months.

If a password reaches the maximum permitted usage period, it must be changed.

For this purpose, the system must automatically inform the user about the expired usage period the next time he logs

on to the system and immediately guide him through a dialog to change the password. Access to the system must no longer be permitted without a successfully completed password change.

For technical user accounts (M2M or Machine-2-Machine), which are used for the authentication and authorization of systems among themselves or by applications on a system, automated solutions must also be implemented to comply with the permitted usage period for passwords.

Alternatively, if such an automatic mapping of the process for changing the password cannot be implemented, an effective organizational measure must be applied instead, which ensures a binding manual password change at the end of the permissible period of use.

Motivation: Unlike more modern authentication attributes, passwords are easier to attack. Without specific measures for reliable, technically automated detection of compromises, the risk of a password being discovered or broken by an attacker can increase considerably over time.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.01-30/9.0

7. Protecting Data and Information

Req 39	Stored data in need of protection must be protected against unauthorized access, modification and deletion.
--------	---

The need for protection of stored data depends on its classification (e.g. according to applicable legal data privacy requirements, regulatory requirements, contractual obligations), the potential damage in the event of its misuse, and other relevant factors (e.g. the location of storage). The nature and extent of protective measures must be appropriately chosen.

Stored authentication attributes such as passwords, private keys, tokens or certificates etc. are generally considered to be in need of protection. Data that determines the functionality and security-relevant behavior of a system (e.g. system configuration files, operating systems and kernels, drivers) are also considered to be fundamentally in need of protection.

Compliance with the protection objectives of confidentiality, integrity and availability must be consistently guaranteed for stored data in need of protection. This also applies during only short-term storage (e.g. when storing in a web cache or in a temporary folder within a data processing chain).

Basically, access to data in need of protection in a system must be fully regulated on the basis of technically implemented authorization assignments and controls.

If such technical access control alone is no longer sufficient to ensure the necessary protection requirements of stored data, or if its effectiveness cannot be consistently ensured, additional cryptographic methods (e.g. encryption, signing, hashing) must be implemented. Cryptographic methods used for the storage of data must be recognised and effective for this purpose and must not have any known vulnerabilities.

Motivation: The storage of data on a system without adequate protection enables an attacker to view, use, disseminate, modify or destroy it without authorization. This potentially opens up additional attack vectors on the immediate and connected other systems and can lead to significant failures, loss of control and damage as well as resulting penalties and loss of reputation towards customers and business partners.

Implementation example: [Example 1]

A system exports data for transport to mobile media. Since the system's technical access control at the file permission level no longer applies as soon as the mobile media is removed from the system, additional measures must be taken to protect the data. Before the system writes the data to the mobile media, it is encrypted accordingly using a suitable algorithm. The associated encryption key is exchanged on a separate channel so that the data can be decrypted and processed again in the legitimate target system. An attacker who takes possession of the mobile media, on the other hand, has no access to the data.

[Example 2]

Only cryptographic hashes of passwords generated with a secure password hashing method are stored in the local user database of a system. For the system, these hashes are sufficient to authenticate users when they log on to the system. However, if an attacker can copy the user database, he does not immediately come into possession of plaintext passwords with which he could log on to the system on behalf of the users.

[Example 3]

On a system, the configuration files of the Web server can only be written by the legitimate admin in which corresponding permissions have been set in the file system. The access control of the operating system kernel thus denies all other users of the system to make changes to the configuration files of the web server; including the web server service account itself, which also reduces the attack surface from the outside in case of vulnerabilities in the web server.

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Unauthorized modification of data
- Disruption of availability
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.01-14/9.0

Req 40	Data in need of protection must be protected against unauthorized access and modification during transmission.
--------	--

The need for protection of data to be transmitted depends on its classification (e.g. according to applicable legal data privacy requirements, regulatory requirements, contractual obligations), the potential damage in the event of its misuse, and other relevant factors (e.g. transmission via public networks). The nature and extent of the protective measures must be appropriately chosen.

Authentication attributes such as passwords or tokens etc. are generally considered to be in need of protection. Data that determines the functionality and security-relevant behavior of a system (e.g. updates & patches, configuration parameters, remote maintenance, control via APIs) are also considered to be fundamentally in need of protection.

Compliance with the protection objectives of confidentiality and integrity must be consistently guaranteed during the transmission of data in need of protection.

As a rule, this requires the implementation of cryptographic methods (e.g. encryption, signatures, Hashes).

Cryptographic methods may

- be applied directly to the data before transmission, which can make subsequent transmission acceptable even via insecure channels
- be used on the transmission channel to create a secure channel and protect any kind of data passing through it
- or be implemented as a combination of both.

Cryptographic methods used for the transmission of data must be recognized and effective for this purpose and must have no known vulnerabilities.

Motivation: The transmission of data without adequate protection enables an attacker to intercept, use, disseminate, modify or remove it from transmission without authorization. This potentially opens up further attack vectors on the immediate target systems as well as connected other systems and can lead to significant failures, loss of control and damage as well as resulting penalty claims and reputational losses towards customers and business partners.

Implementation example: [Example 1]

Confidential documents are encrypted before they are sent by e-mail to the customer.

[Example 2]

An administrator configures a new cloud application over the Internet. Access is via a TLS-encrypted connection ("https").

[Example 3]

A system obtains automatic software updates from an update server. The update server delivers the software updates cryptographically signed. The system can thus validate the received software updates and reliably rule out that they have been manipulated during transmission.

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Unauthorized modification of data
- Disruption of availability
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.01-15/9.0

Req 41	Key Management System for secure storage of encryption keys must be used. For confidential data, the key management system must be under DTAG's full control.
--------	---

Attacker who manages to obtain the data from the cloud will most likely deal with encrypted data, but the data would be useless without encryption keys (or sufficiently advanced quantum computers). Therefore, it is important to store all the encryption keys in appropriately secure key management system (KMS). The term encryption key here refers to the

key used for symmetric encryption (where encryption key is the same as decryption key and which is the most common use case for data storage in the cloud), or the private key from the public-private keypair in case of asymmetric encryption; or in other words - it is about protecting the key which would enable attacker to convert encrypted data into clear text.

KMS might come in various flavors:

- under full control of CSP (mostly transparent to customers)
- provided by CSP, but keys managed by customers
- DTAG's own KMS collocated at CSP
- DTAG's own KMS

For open and internal data, use of KMS provided by CSP is generally allowed. The last solution (DTAG's own KMS) is the only acceptable solution and allowed for confidential data since it protects the data also against a rogue cloud provider or intelligence agencies.

Not all systems and data hosted in the cloud will have same secrecy requirements, so it is necessary to select right option for the given system or data. The BoM decision typically includes the level of data which can be hosted in the cloud (and under which conditions), and it would be possible to have a cloud provider approved (or some services of the cloud provider) which, despite technical capabilities, must not be used for storing confidential (or personal) data.

Existing security requirements for KMS still apply, or in other words - use of KMS for the cloud does not change requirements for key strength etc. Also, communication with KMS needs to be encrypted.

Most SaaS cloud services do not allow or support using any kind of external key management, but for those which do, an appropriate solution has to be used according to the data classification.

This requirement is in line with Requirement 16 from 3.63. ("Operations in DT IT") for those systems where 3.63. applies.

Motivation: A KMS enables fine-grained access control and auditability of key usage, providing a robust security foundation for protecting data with need of protection at rest and in transit. KMSs usually have HSMs on the backend that provide physical protection, secure key storage, and encryption acceleration, helping to ensure the confidentiality, integrity, and availability of critical cryptographic operations.

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Unauthorized modification of data
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-52/5.0

Req 42 Data lifecycle must be configured through available policies.

Public cloud providers typically offer configuration of data lifecycle through policies, which is most typically used for object storage. Definition of data lifecycle allows automated movement of data through different storage tiers and appropriate restriction levels until data is deleted. Automation minimizes risk of data leaks and improves regulatory compliance. Main benefits include:

- moving data from hot to cold storage tiers (from regularly used data to archived data), which is typically done for economic reasons
- minimal lifetime of the data necessary by regulatory requirements
- maximal lifetime of the data allowed by regulatory requirements
- automated deletion of the data at the end of lifecycle
- support for setting up more complex lifecycle which includes data operations, e.g. anonymization or tokenization of data at certain point in time

Keys used to encrypt and decrypt data need to follow the same lifecycle.

Motivation: Data lifecycle streamlines data management, reduce storage costs, comply with regulatory requirements, protect information with need of protection, and maintain data integrity throughout the defined process of the lifecycle.

For this requirement the following threats are relevant:

- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.66-55/5.0

8. Protecting Availability and Integrity

Req 43	Systems which need high availability and reliability must be designed and deployed in appropriately redundant way.
--------	--

Cloud service providers have typically very resilient infrastructure, but that does not apply to every single component in the cloud. Virtual machines and containers die, even more frequently compared to traditionally built counterparts. Systems which need high availability and reliability in the cloud need to adapt their HA mechanisms, so they remain highly available and resilient in cloud environments, i.e. systems and applications need to be designed to handle failures. Using different regions, availability zones or data centers of CSP (or even multiple cloud providers) for redundant application components is preferred choice for ensuring HA of the overall system deployed in the cloud.

One could consider using chaos engineering to ensure that a system deployed in the cloud remains available and resilient during failures (e.g. Chaos Monkey or Simian Army).

Most SaaS cloud services do not allow or support configuring specific regions or simply have no different tiers, but for those services which do, an appropriate region or tier has to be used to achieve desired SLA and satisfy high availability, business continuity and disaster recovery expectations.

Motivation: Redundancy helps mitigate the impact of hardware failures, software glitches, or natural disasters, ensuring uninterrupted service delivery, reducing downtime, and enhancing overall system reliability and performance.

For this requirement the following threats are relevant:

- Disruption of availability

ID: 3.66-59/5.0

Req 44	Backup solution must be used. And restore function must work.
--------	---

The requirement can be skipped if there is no need to preserve the data.

Cloud providers typically offer excellent reliability, availability and durability of the data stored in the cloud, except for cases when their data centers burn down (which has happened). They, on the other hand, do not perform backup for customers automatically and it is responsibility of the system owner to implement a backup and restore strategy according to business needs for the given system, with considering the following scenarios:

- accidental or intentional (insider) data overwrite/deletion
- data corruption due to bugs
- ransomware attack
- fire, flood, earthquake or any other disaster which can impact a data center, availability zone or even a region of the cloud provider

It is worth mentioning that mounting a backup disk on a system which is being backed up is probably worst practice in case of ransomware attack because ransomware would lock both the system and the backups. If possible, a backup solution should use a backup storage type which allows writing data once, but no further changes of already written data and with policies which prevent too early deletion of backups. Also, backup system (e.g. backup disk or storage bucket) should generally belong to another tenant with different credentials so that an accident, attacker or malware cannot destroy the backups at the same time as the system (one may think of an analogy to running `'rm -rf /'` with backup disk mounted).

Backup data needs to include all data necessary for business continuity, e.g. "regular data", system configurations, encryption keys (which must be handled separately), logs, ...

The backups need to be protected against unauthorized access, typically by using encryption and protecting the encryption keys accordingly.

Backup without working restore function is not enough, so the overall concept needs to include restore mechanism

and restore tests, which are preferably automated. Testing procedure for the restore must be designed in a way that there is no possibility to accidentally cause downtime or data loss to the production system or cause any other kind of data leak.

Depending on the data classification, backups can reside:

- in the same cloud (from same CSP)
- in another cloud (from another CSP)
- in DTAG own premises (would be unusual approach, but still possible)

The prioritization, or order, of systems which need to be restored from the backup needs to be established upfront since this is generally a time consuming process and performing the prioritization during an incident would jeopardize service level objective (SLO) and recovery time objective (RTO). Further details can be found in requirement 5 in document 3.61 Operations.

Although SaaS cloud services do not offer many capabilities for backup, they may still offer versioning, prevention of (mostly accidental) deletion, and other features for preservation of data. Such features have to be appropriately configured to achieved desired data durability.

Motivation: Backups ensure data resilience, mitigate the risk of data loss, and facilitate data recovery in the event of system failures, disasters, or security incidents. Backups serve as a safeguard against accidental deletions, hardware failures, ransomware attacks, or other unexpected events that can result in data corruption or loss.

For this requirement the following threats are relevant:

- Disruption of availability

ID: 3.66-60/5.0

9. Architecture and Networking

Req 45	Each application/system including its environment type (e.g. development, test, production) in the cloud must be separated from any other applications/systems and environments in form of a tenant/project, including separation between the application/system and respective backup, and each tenant/project must be handled in the same way regardless of the environment type.
--------	---

Applications must be (logically) separate from others, in a way of separate tenants (projects). Separation is also needed for different environment type so that there is no accidental transfer of data between environment or no unwanted actions (e.g. preventing leaking test data to public, preventing leaking production data to development environment, or preventing accidental deployments of software or configuration to production environment instead to development environment).

Separation of application and backup of the same application, by using another tenant and separate credentials, is also needed to preserve backup in case the application is compromised.

The separation also applies to the management part of the cloud application/system. Example could be element managers or CI runners which must be separate per environment type, despite development, test and production part of the application using same (multitenant, of course) CI/CD chain.

This approach enables mandatory segregation of duties and prevents unwanted access to data and systems, while also allows to control communication between different services.

System or application is a very broad term and it may refer to anything between a single microservice to a whole system providing IPTV platform. In this context, it is left to the system owner to define together with security experts the most reasonable and balanced separation, but the extremely rough guidelines are:

- a single microservice does not require its own tenant/project
- huge systems like IPTV might need to be broken down into several tenants based on functional building blocks (e.g. provisioning and management engine, live stream, video library)

In addition to separation of accounts, it is necessary to protect each project/tenant in the cloud in the same way. It is understandable that tenants of production systems will have the highest possible protection, but development and test tenants also deserve the same level of security. The credentials of all tenants need to be protected with due care, they must not end in Git repositories in clear text and they can be only used from approved workstations (corporate managed devices or development workstations). It is also recommended that development and test environments have lower quotas and earlier alarming to detect the misuse in timely manner. This prevents one of most common types of public cloud abuse - cryptomining, most often caused by leaked credentials due to human mistake or use of insecure development workstations.

This requirement is in line with requirements from 3.14. ("Architecture of systems") and applies in addition to them in case both requirement documents are used together.

Motivation: In the case of test and development systems, it cannot be assumed that the system state is sufficiently safe, especially since they are inherently exposed to permanent changes. If the system types are linked to each other or the same tenants/projects are used, there is a risk that unauthorized persons can access production systems and impact data from the test/development environment or that the stability of the production systems and thus the availability of the associated services will be jeopardized. A test and development infrastructure that is completely independent of production makes it possible to carry out changes for test and development systems even during a "frozen zone". This allows security updates to be tested and deployed to production more quickly.

Implementation example: For SaaS services, the separation can be achieved by using different accounts, folders or sections, depending on the terms used by the SaaS cloud provider.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized use of services or resources

9.1. Cloud Connectivity

Req 46 The accessibility of activated services must be restricted.

In principle, a service provided must be completely deactivated on all interfaces of the system through which accessibility of the service is not required for the proper operation of the system. The deactivation is primarily to be implemented by a corresponding configuration of the service or operating system. In cases where the available configuration options do not allow deactivation on individual interfaces, a local filter ("Host Firewall") may instead be used on the system to block access to the service via unnecessary interfaces.

The accessibility of a service via the required interfaces must also be restricted to legitimate communication partners. The restriction must be implemented by a corresponding configuration of the service or operating system or by means of a local filter ("Host Firewall"). Alternatively, this task may be outsourced to a network-side filter element, provided that the system is located in a separate network segment and communication with this segment is only possible via the network-side filter element.

Motivation: By deactivating services on interfaces through which accessibility is not necessary, as well as by restricting possible communication partners, the attack surface offered by a system can be greatly reduced.

Implementation example: An SNMP service used to monitor a system is enabled exclusively on the dedicated management network interface of the system. A firewall also regulates that only the legitimate monitoring system of the infrastructure environment can reach this service.

For this requirement the following threats are relevant:

- Unauthorized use of services or resources
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

ID: 3.01-6/9.0

Req 47 Management services of applications/services in the cloud must not be widely reachable.

SSH and RDP are most common remote management services and they play an important role in securing cloud applications. An attacker who manages to get access to management services can compromise the whole system or even more systems with very low effort. Limiting the attack surface is an important step to mitigate possibility of such attacks. Management services of an application in the cloud must not be widely reachable. They should not be reachable over Internet, however, if connectivity over Internet is necessary, the access must be restricted to trusted parties and/or centrally managed and monitored jump servers need to be used as an entry point.

Use of immutable images typically requires no management access, therefore, it reduces attack surface.

If SaaS cloud provider offers such capability, the access to the administration portal can be restricted to certain geographical locations, IP addresses, specific days of the week or specific times of the day. By doing this, one should consider not to completely lock out administrative users, especially so called accounts used for emergency access.

Motivation: Making common remote management services widely accessible exposes them to potential threats from malicious actors attempting to gain unauthorized access to systems or launch password-guessing attacks. Limiting the attack surface is an important step to mitigate the possibility of such attacks.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources

ID: 3.66-66/5.0

9.2. Cloud Connectivity - External Connectivity

Traffic going out of the cloud or coming in to our premise within the cloud.

Req 48	Connections between public clouds and resources and networks inside DTAG premises must be protected by an independent network filter element, located inside DTAG premises and under DTAG control. A service running in the cloud must not have direct access to Internet and internal DTAG networks at the same time.
--------	--

Connections between public cloud and internal networks of DTAG (e.g. CN-DTAG, Intranet) must be protected with a security device which needs to, depending on specific use case, provide:

- stateful firewalling (as bare minimum)
- web application firewall
- application level gateway
- API gateway
- DoS/DDoS protection
- ...

The security device in this context may contain out of several physical and virtual devices acting as a single unit. Protection provided by the device must be in general bidirectional - mainly protecting DTAG's networks from threats from the public cloud, but also protecting resources inside public cloud from malware and malicious insiders inside DTAG's networks. Since it is expected that applications hosted in the public cloud will utilize security groups and other security features offered by the public cloud itself, controls necessary on the security device do not to be very detailed for the connectivity from DTAG internal network towards the network in the public cloud. For the other direction, protection of DTAG's internal networks from threats from the public cloud, implementation must be done fully on the security device.

For HTTP(S) based services, the provided device should protect against OWASP Top 10 web vulnerabilities at minimum.

The configuration of the device should be automated in a way so that it can cope with rapid provisioning cycles typically found in highly automated cloud environments (e.g. it should be handled so that it can cope with dynamic IP address assignments in cloud).

Internal networks of DTAG are not built in modern way with microsegmentation in mind, and applications deployed in the cloud need to follow guidelines for connectivity to existing DTAG internal networks if they need such connectivity. Applications in the cloud are not allowed to have direct Internet connections and connection to DTAG internal networks at the same time because there would be no (at least) 2 layers of security protection in that case. This would allow attackers much easier infiltration into DTAG network and resources through a backdoor placed, intentionally or unintentionally, in an application running in the cloud. For special use cases, e.g. management and monitoring of an Internet facing application running in the cloud, an appropriate solution has to be determined on case-by-case basis.

A service running in the cloud which has access to both Internet and internal DTAG networks must have the connectivity split over multiple VMs, containers or in general - multiple tiers.

Motivation: DTAG's internal networks are not built in a modern way with microsegmentation in mind, and applications deployed in the cloud must follow guidelines for connectivity to existing DTAG internal networks if they need such connectivity. Applications in the cloud must not have direct internet connections and connections to DTAG's internal networks at the same time, as in this case there would be no (at least) 2 layers of security. This would allow attackers to penetrate the DTAG network and resources much more easily through a backdoor that is intentionally or unintentionally placed in an application running in the cloud.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized use of services or resources
- Unnoticeable feasible attacks

Req 49	DTAG internal procedures must be defined that allow to cut off any connections between the CSP and DTAG own infrastructure.
--------	---

Malware, especially ransomware, can spread very fast via network and can also impact data available via network even on storage systems which require authentication. In certain cases, the best action is to block network communication. This applies not only to internal systems, but also to cloud environments.

Motivation: In case of incidents like ransomware attack it must be possible to cut off cloud environments immediately to avoid further spreading.

Implementation example: The procedure defined for cut off scenario needs to consider necessary steps to cut off communication by either unplugging cables, disabling network ports, shutting down VPNs, blocking all traffic on firewall or any other available method depending on the connection type. The scenario needs to consider that the ransomware infection may start in the cloud environment requiring isolation of the cloud environment, but also a possibility that the infection started in DTAG network and is spreading toward the cloud environment.

For the cases where connectivity between the cloud and DTAG is only via Internet (like SaaS), the blocking of connectivity may be performed also by blocking affected accounts.

For this requirement the following threats are relevant:

- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Unauthorized use of services or resources
- Disruption of availability
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

10. Logging and Monitoring

Req 50	Applicable retention and deletion periods must be observed for security-relevant logging data that is recorded locally.
--------	---

From an IT security perspective, local storage of security-relevant logging data on a system is not mandatory. Since the local storage can be damaged in the event of system malfunctions or manipulated by a successful attacker, it can only be used to a limited extent for security-related or forensic analyses. Accordingly, it is relevant for IT security that logging data is forwarded to a separate log server.

Local storage can nevertheless take place; for example, if local storage is initially indispensable when generating the logging data due to technical processes or if there are justified operational interests in also keeping logging data available locally.

The following basic rules must be taken into account when storing logging data locally:

- Security-related logging data must be retained for a period of 90 days.
(This requirement only applies if no additional forwarding to a separate log server is implemented on the system and the logging data is therefore only recorded locally.)
- After 90 days, stored logging data must be deleted immediately.

Deviances

Different retention periods and deletion periods may exist due to legal or regulatory requirements (especially in connection with personal data) or may be defined by contractual agreements. In these cases, the applicable periods must be agreed individually with a Project Security Manager (PSM) / Data Privacy Advisor (DPA) or are specified by them.

Motivation: Logging data is an immensely important IT security tool for preventing, detecting and clearing up system faults, security and data privacy incidents. On the other hand, the recording of logging data, like any other data processing, is also subject to legal and regulatory requirements. Accordingly, guidelines must be adhered to that reconcile the two.

Implementation example: Taking into account the current legal situation and applicable data privacy regulations, the following deletion periods for locally stored security-relevant logging data are implemented on an exemplary telecommunications system:

- Standard System Logs: Deletion after 90 days at the latest
- Logging of public IP addresses: Deletion (or anonymization) after 7 days at the latest
- Logging of the assignment of dynamic public IP addresses by the telecommunication solution: Deletion after 7 days at the latest
- Logging of non-billing-relevant call detail records: Deletion after 7 days at the latest
- Logging of the content of e-mail and SMS: Deletion after 24 hours at the latest
- Logging of the domain queries handled by the DNS server of the telecommunications solution: Deletion after 24 hours at the latest

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Denial of executed activities
- Unnoticeable feasible attacks

ID: 3.01-34/9.0

Req 51	Security-relevant logging data must be forwarded to a log server immediately after it has been generated.
--------	---

The system's security-relevant logging data must be forwarded to a log server immediately after it has been generated. For the purposes of the requirement, the log server must meet the following criteria:

- The log server must be dedicated to the storage of logging data. It must not depict any other functions or services that are not related to this task. A permissible exception is the combination of a log server with a SIEM (Security Information and Event Management) solution.
- The log server must be separate from the system. It must not be operated directly on the system as an additional module.
- The log server must be centrally available for the operational environment, to which the system belongs.

To transmit the logging data to the log server, standard protocols specialized for this purpose (e.g. Syslog or SNMPv3) should preferably be used

Motivation: If logging data is stored only locally, it can be manipulated in the event of a successful compromise of the system by an attacker in order to obfuscate their attack and actions taken on the system. For this reason, the forwarding of logging data to a log server must take place very soon after the occurrence of a security-relevant event in the source system. This means that the logging data generated is separated from the source system and is available for analysis at any time, regardless of the source system state. The use of a central log server for an operational environment ensures that the logs of all IT/NT Systems in this operational environment can be accessed quickly and uniformly and that the access itself can be optimally integrated into the incident processes of the environment.

For this requirement the following threats are relevant:

- Unauthorized modification of data
- Disruption of availability
- Denial of executed activities
- Unnoticeable feasible attacks

ID: 3.01-35/9.0

Req 52	For security-relevant logging data that is forwarded to the log server, compliance with the applicable retention and deletion periods must be ensured.
--------	--

The following basic rules must be taken into account:

- security-related logging data must be retained for a period of 90 days on the log server.
- after 90 days, stored logging data must be deleted immediately on the log server.

Deviances

Different retention periods and deletion periods may exist due to legal or regulatory requirements (especially in connection with personal data) or may be defined by contractual agreements. In these cases, the applicable periods must be agreed individually with a Project Security Manager (PSM) / Data Privacy Advisor (DSB) or are specified by them.

Log server under the responsibility of a third party

If the selected log server is not within the same operational responsibility as the source system of the logging data, it must be ensured that the responsible operator of the log server is aware of the valid parameters for the logging data to be received and that they are adhered to in accordance with the regulations mentioned here.

Motivation: Logging data is an immensely important IT security tool for preventing, detecting and clearing up system faults, security and data privacy incidents. On the other hand, the recording of logging data, like any other data processing, is also subject to legal and regulatory requirements. Accordingly, guidelines must be adhered to that reconcile the two.

Implementation example: Taking into account the current legal situation and applicable data privacy regulations, the

following deletion periods for forwarded security-relevant logging data from an exemplary telecommunications system are implemented on the separate log server:

- Standard System Logs: Deletion after 90 days at the latest
- Logging of public IP addresses: Deletion (or anonymization) after 7 days at the latest
- Logging of the assignment of dynamic public IP addresses by the telecommunication solution: Deletion after 7 days at the latest
- Logging of non-billing-relevant call detail records: Deletion after 7 days at the latest
- Logging of the content of e-mail and SMS: Deletion after 24 hours at the latest
- Logging of the domain queries handled by the DNS server of the telecommunications solution: Deletion after 24 hours at the latest

For this requirement the following threats are relevant:

- Unauthorized access or tapping of data
- Denial of executed activities
- Unnoticeable feasible attacks

ID: 3.01-36/9.0

Req 53	Security-relevant logging data must be provided to a security incident detection and response solution in near-real-time.
--------	---

An acceptable detection and response solution for the purposes of this requirement must meet the following criteria:

- There is a continuous analysis and correlation of the logging data provided by connected IT/NT Systems.
- The identification of threatening activities and ongoing attacks takes place on the basis of a continuously updated detection pattern database
- Integrated alerting processes trigger the initiation of countermeasures in the event of a find.
- Security experts are actively involved in the operation of the solution as well as the evaluation of alarms

Such a solution is not the same as a log server in terms of its function. Accordingly, the provision to a detection and response solution must always be in addition to an existing forwarding to a log server. An exception exists if the log server used provides the logging data to a detection and response solution immediately after receipt. In this case, it is sufficient for the system to forward the logging data only to the log server.

The logging data must be provided to the detection and response solution within 5 minutes ("near-real-time") from the time it is generated in the system.

The specific technical interfaces and protocols to be supported for deployment are determined by the detection and response solution available in the system's operating environment as well as the technical basis of the system itself.

Motivation: IT/NT systems are constantly exposed to security threats and attacks. By means of a central analysis and the correlation of recorded system events of entire IT landscapes in almost real time, current threats and ongoing attacks can be detected even beyond system boundaries and appropriate countermeasures can be initiated promptly. The prompt delivery of logging data to the detection and response solution is essential for its effectiveness.

Implementation example: **Example 1:** A system based on a classic server collects security-relevant logging data in the operating system's local syslog or Windows Event Log. Logging data is continuously forwarded to a Security Information and Event Management (SIEM) for further analysis, intrusion detection, and alarm handling.

Example 2: An Endpoint Detection and Response (EDR) agent is running on a system. The EDR agent continuously forwards the security-related logging data it collects on the system to the central backend of the EDR solution, where further analysis, intrusion detection, and alarm handling takes place

Example 3: Network components provide events to a SIEM that make it possible to detect suspicious communication behavior such as network scans, unauthorized data flows between network segments, or active connections or connection attempts to suspicious destinations on the Internet. In particular, gateway routers from the CNDTAG provide

"Netflow" logs, while forward proxies report access to domains and/or IP addresses and firewalls log blocked connection attempts.

Example 4: Systems designed on the basis of cloud-native functions generate logging data directly from these functions using the cloud's own logging and monitoring infrastructure. In such a system, the generation of security-relevant logging data is specifically activated in the corresponding configuration options of the cloud-native functions used. The resulting logging data is continuously transferred to a SIEM (Security Information and Event Management) for further analysis, intrusion detection and alarm handling. Some cloud providers already provide some attack pattern detection in the cloud's own monitoring infrastructure and generate specific cloud alerts from it. These are also transferred to the SIEM.

Example 5: A system is based on a SaaS ("Software as a Service") application from a cloud provider. Typically, in the SaaS model, the cloud provider is responsible for all operational aspects of the application, including security management. The cloud customer itself has no influence on the underlying technical infrastructure or the operation of the application. Based on contractual agreements, recognized certificates, and documented self-disclosures from the cloud provider, it is formally verified that the cloud provider has implemented a security incident detection and response solution. Some cloud providers also generate dedicated logging data about security-relevant activities within the separate customer tenants of their SaaS application (e.g. changes to configuration settings, permission assignments or user accounts within the tenants) and make them available to cloud customers. These so-called audit logs are continuously transferred to a SIEM (Security Information and Event Management) for further analysis, attack detection and alarm handling.

For this requirement the following threats are relevant:

- Unnoticeable feasible attacks

ID: 3.01-37/9.0

Req 54	Security relevant events in the cloud must be logged and forwarded to a separate log storage and to the Security Information and Event Management System (SIEM) of the responsible SOC to allow 24/7 security monitoring.
--------	---

General security requirements on logging still apply and are also relevant for cloud environments. Security relevant events like cloud alerts, cloud management and API call logs, configuration changes, access to accounts, access to certain data with need of protection, network logs, system logs etc., must be logged in proper way by using logging and monitoring infrastructure of CSP, for all service types used - IaaS, PaaS and SaaS. Supportive systems like IAM and KMS used for the cloud must be also included in the logging concept. And lastly - existing other security requirements specify that all systems, including those deployed in the cloud, like operating systems running in virtual machines, databases, server software and containers, must provide logs.

Logging and forwarding of security relevant events must be configured by the landing zone for all tenants/accounts/systems in a way that it happens immediately after generation and cannot be de-activated by admins/users of the tenants/accounts/systems. If this cannot be ensured by the landing zone, the landing zone has to offer solutions that the tenants/accounts/systems must use (e.g. by documentation on configuration of log forwarding, use of pre-configured images) and established controls that forwarding is enabled.

Monitoring of security relevant events is a prerequisite for detecting attacks on cloud resources like data exfiltration or resource consumption. The monitoring must be 24/7 for all environments (including test and development environments), because in addition to traditional attacks with malware and ransomware on production environments, cryptocurrency mining by attackers in development environments over the weekend is also a threat which can cause huge financial damage.

Forwarding of logs to both a separate log storage and to SIEM at the same time applies to those cases where SIEM does not store logs for necessary period of time, like it is the case for CDC (Telekom Security Cyber Defense Center - CDC is commonly used in Germany) who perform security monitoring for DT IT. In that specific case, the separate log server stores the logs for forensic analysis and to have evidence in legal actions against attackers.

Logging policy must follow laws, regulations, and company policies. According to current laws, regulations and company policies, logging of such events is permitted only for defined use cases and forbidden for others, like work control of employees.

Examples of some event types/logs:

- cloud alerts: GuardDuty for AWS or SecurityCenter in Azure
- cloud management and API call logs: AWS CloudTrail or Azure Activity Logs
- network logs: all connections to and from the Internet must be logged by proxy and if feasible all relevant network traffic - VPCFlowlogs
- system logs: syslog/auditd and winlog

In case of a system where 3.63. ("Operations in DT IT") applies, Requirement 25 from 3.63. applies as more specific requirement.

Motivation: If logging data is only stored locally, it can be manipulated by an attacker who succeeds in compromising the system in order to conceal his attack and any manipulation he has performed on the system. This is the reason why the forwarding must be done immediately after the event occurred.

Only if log data are analyzed, threats can be identified. Analyzing needs to be done 24/7, because attackers do not stop their attacks outside business hours.

For this requirement the following threats are relevant:

- Denial of executed activities
- Unnoticeable feasible attacks

ID: 3.66-80/5.0

Req 55	Incident response process for reacting appropriately to security events and incidents must be established, including provisioning of forensic evidence.
--------	---

Clearly defined incident response process is mandatory to minimize the impact of an ongoing or successful attack. The process defines responsibilities, contact persons and tasks for attack mitigation, analysis of potential or real damage, procedures to recover system security and preserve evidence for law enforcement.

Consideration of both sides - a system deployed in the cloud and APIs available from the cloud provider is necessary for an incident response process, e.g. in case of dealing with hacked virtual machines it may be beneficial to create snapshots of hacked VMs through APIs of the cloud provider, so they can be preserved as an evidence.

In alignment with the responsible SOC, the landing zone operators must provide incident response capabilities towards SOC with appropriately short reaction time and 24/7 availability in order to be able to directly stop attacks and prevent further damages. The capabilities must cover all environments (including test and development environments) and access to all tenants/accounts/resources must be possible even without system owners (e.g. via break glass accounts). Processes must be defined, established, and regularly tested that allow to get always back control over compromised resources and to properly collect the forensic evidence. Examples include:

- Isolation of compromised virtual machines
- Update other relevant security groups
- Revoke misused users, roles or rights and
- Stop malicious services or serverless functions
- Export further relevant data such as copies of virtual machine images, containers and storage to the forensic laboratory of responsible SOC

In case when the attack was not recognized by SOC but the system owner, responsible SOC must be informed about security events and incidents and it is expected that the responsible SOC will provide enough information on how to proceed with the investigation.

An illustration of a few good and bad practices is listed below.

Examples for proper forensics:

- Creating a snapshot of infected containers and/or virtual machines with all relevant logs.

- Isolate impacted workloads and keep them fully isolated (full network and storage separation from rest of healthy environment) and wait for further instructions from a security officer.

Examples of bad forensics or destroying forensics data, such actions must NEVER happen:

- Deleting containers and/or virtual machines because everything is already logged.
- Keep everything running "normally" (as is) when it gets infected/compromised.

In case this requirement is fulfilled, Requirement 8 in 3.63. ("Operations in DT IT") can be skipped.

Motivation: Security incidents must be dealt with quickly and competently in order to detect and avert potential damage to the company, employees and customers.

For this requirement the following threats are relevant:

- Denial of executed activities
- Unnoticeable feasible attacks

ID: 3.66-81/5.0

11. References

Some references from the document:

<https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>

<https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-144.pdf>

<https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-86.pdf>

12. Glossary

Public, private, community and hybrid -cloud definitions are taken from NIST.

Public cloud

The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

Private cloud

The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g. business units). It may be owned, managed, and operated by the organization, a third party or some combination of them, and it may exist on or off premises.

Community cloud

The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

Hybrid cloud

The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g. cloud bursting for load balancing between cloud).

IaaS, PaaS and SaaS definitions are taken from NIST.

IaaS

The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g. host firewalls).

PaaS

The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

SaaS

The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface.

The consumer does NOT manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities; with the possible exception of limited user-specific application configuration settings.

Additional definitions

CSP stands for Cloud Service Provider. Examples are Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP).

Landing zone refers to a configuration of cloud environment, typically fulfilling security and compliance aspects, identity and access management, networking, billing etc.. In context of this document, the term landing zone is often also called administration portal or administration cloud console.

Personal user account is a user account used by a person (human being).

Technical user account is a user account used by a machine, also known as M2M (machine to machine) user account, or simply as service account. It is mainly used in M2M communication scenarios where communication is established between two machines, typically without human interaction.

VPC stands for "Virtual Private Cloud," and it is a virtual network environment within a public cloud service (e.g., AWS, Azure, GCP) that allows users to provision and manage their cloud resources in an isolated and private network.

VNET stands for "Virtual Network," and it is a virtualized network infrastructure within the Azure cloud platform. VNET allows users to create and manage their isolated network environments, define IP address ranges, and control traffic flow between Azure resources and the internet or on-premises networks, ensuring secure and efficient communication within the Azure cloud environment.