Security requirement

Architecture of datacenter and cloud infrastructures

Deutsche Telekom Group

Version	2.4
Date	Jul 1, 2023
Status	Released

Publication Details

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

Reuterstrasse 65, 53315 Bonn
Germany

File name	Document number	Document type
	3.58	Security requirement

Version	State	Status
2.4	Jul 1, 2023	Released

Contact	Validity	Released by
Telekom Security psa.telekom.de	Jul 1, 2023 - Jun 30, 2028	Stefan Pütz, Leiter SEC-T-TST

Summary
Security requirements for the architecture of datacenter and cloud infrastructures.

Table of Contents

1. Introduction

This 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 in 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.

2. Infrastructure

Req 1	A datacenter infrastructure or cloud platform must be able to separate systems and data, operated on top of this platform, in accordance to their protection needs.

Critical systems, in particular, always have to be separated from other systems including other critical systems. Similarly, systems on which personal data is processed must be protected against unwanted access or data flows from the same network or other networks.

Motivation: It is more likely for a less protected system to be compromised. This must not result in other systems with higher protection requirements being attacked by this compromised system.

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
- Disruption of availability
- Unnoticeable feasible attacks

For this requirement the following warranty objectives are relevant:

ID: 3.58-1/2.4

Req 2	If systems are accessible from external networks (e.g., Non-DTAG networks such as the internet), they must be implemented on an infrastructure which is physically separated from internal systems.

To support the required separation in the layer model, all externally reachable machines and the corresponding infrastructure elements like hypervisors, and network elements such as switches, routers, firewalls and loadbalancers must be physically separated from internal systems.

Motivation: This requirement ensures that there is only one path from the systems accessible from the Internet to the underlying application and database systems. The firewall or filter elements cannot be circumvented via components that are accessible from the Internet as well as from internal networks. This has also the benefit that infrastructure incidents only affect a sub-part while other parts remain available.

Implementation example: Independent switches and routers are used to connect a presentation layer to the internet.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Disruption of availability
- Unnoticeable feasible attacks

For this requirement the following warranty objectives are relevant:

ID: 3.58-2/2.4

Req 3	If infrastructure elements provide services or functions for internal and also for externally accessible machines, these functions must be separated with respect to internal and external aspects.

Infrastructure elements includes network components such as routers and switches, common used services like DNS or NTP, management and monitoring systems as well as central storage systems. For these systems it must be ensured that internals are not accessible from the outside.

Motivation: Internal data must not be compromised by external reachable components. Implementation of need-to-know and need-to-see principles.

Implementation example: An implementaion example can be separated usage of storage media, separate SAN fabrics or virtualization functions (storage virtualization) like vFiler.

For this requirement the following threats are relevant:
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Disruption of availability

For this requirement the following warranty objectives are relevant:

ID: 3.58-3/2.4

Req 4	Common used infrastructure services must especially be secured and available.

Failure of main services and security equipment can have an adverse effect on the protection of systems. They must be designed with the relevant security and availability in mind and not have any single point of failures. This includes:

- DNS servers for name and address resolution
- NTP servers for time synchronization
- Outgoing (web) proxies for controlled internet access
- Jumphosts for system access internally as well as 3rd party access
- System management and monitoring
- Installation servers, Patch and Update servers, Software Repositories

Corresponding redundancies also make it possible to maintain central systems in ongoing operations. If the infrastructure extends across multiple sites, it is advisable to make these central functions also available across these locations.

Motivation: Main infrastructure services have a high need of protection and must especially be secured and available.

Implementation example: Jumphosts should be redundant, so that all systems can be administrated all the time, even by third parties. Outgoing (web) proxies should always be able to deliver software updates.

For this requirement the following threats are relevant:
- Unauthorized use of services or resources
- Disruption of availability

For this requirement the following warranty objectives are relevant:

ID: 3.58-4/2.4

Req 5	IP addresses of infrastructure elements must be configured statically.

Motivation: Avoidance of unintended change of IP addresses of these elements.

For this requirement the following threats are relevant:
- Disruption of availability

For this requirement the following warranty objectives are relevant:

ID: 3.58-5/2.4

Req 6	Attacks on the infrastructure platform must be recognized to ensure early countermeasures.

In order to detect attacks on the platform as early as possible and to minimize potential impacts, it is necessary to implement appropriate measures for monitoring and correlation of security events.

Motivation: Especially virtual environments provide additional layer of attack surface by using virtualisation software. This must be monitored and protected through the appropriate safety features, such as those already used in other technology areas. Only through active monitoring, detection and resulting subsequent actions, a continuous protection is possible on operated machinery.

For this requirement the following threats are relevant:
- Disruption of availability
- Unnoticeable feasible attacks

For this requirement the following warranty objectives are relevant:

ID: 3.58-6/2.4

3. Virtualization

Req 7	Data and functions of virtual machines must be reliably separated by the virtualization environment.

Virtual machines, including their data, services and functions offered in the network, must be protected from each other. To achieve this, security functions of the virtualization solution must be used across all levels (physical, logical, network based, in the management, etc.).

Especially it has be ensured, that a compromised virtual instance does not allow access to another virtual instance or to the hypervisor. That means virtual machines must be protected against each other and no direct communication must be possible.

Virtual machines and their data must be grouped in such a form, that in the case of a break of the security features of the virtualization environment, the potential damage is limited to an acceptable level (residual risk). A consideration of the extent of possible damage is just as necessary as an assessment of protection needs, criticality, of the threat potential of individual systems, as well as the data to be processed. These indicators are to assess and to consider in the planning and implementation.

Motivation: The consistent partition of virtual instances limits the risk to compromise all virtual instances. To keep the potential effects of an attack as low as possible, for example in the form of an acquired virtual machine to another virtual machine, including the data, is to implement additional segmentation.

Implementation example: A corresponding grouping could, for example, follow the n-tier architecture model. By implementing, compromised virtual machines that are placed in the presentation layer, a vulnerability in the virtualization layer, granting no access via the hypervisor to the application and database layers on other physical machines.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Disruption of availability
- Unnoticeable feasible attacks

For this requirement the following warranty objectives are relevant:

ID: 3.58-7/2.4

Req 8	Deployment of virtualization must not make it possible to circumvent security interfaces, functions and intrusion detection systems.

Using virtualization may sometimes make it possible to circumvent certain specified security mechanisms, e.g., the authorizations, rules or routes configured on a host or the underlying network. Circumvention of these security functions may jeopardize the security of other systems.

Motivation: If network gateways are interconnected directly via the virtualization software, attacks cannot be detected, i.e., network services are not protected from unauthorized access.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Unauthorized use of services or resources

For this requirement the following warranty objectives are relevant:

ID: 3.58-8/2.4

Req 9	The networks needed for the operation of a virtualization environment (hypervisor management, live migration, heart beat, ...) must be completely separated from the networks (including management network) of the virtual machines running on top of the virtualization environment.

Traffic data on the virtual machines must be separated from the administration of the virtualization environment. A physical separation is recommended.

Motivation: The (hardware based) platform and the corresponding management must always be under control of the operator. It is designed to prevent a virtual machine influencing the virtualization environment or gaining access to it.

Implementation example: Use of one or more physical interfaces to manage the hypervisor, possibly VLANs, and one or more additional interfaces for the networks of the VMs.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Disruption of availability
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

For this requirement the following warranty objectives are relevant:

ID: 3.58-9/2.4

4. Networks

Req 10	Production networks (external and internal), management networks, office networks, test & development networks, transport networks and other network areas must be separated from each other.

A physical separation of these network areas is to be preferred.

If there are not enough physical adapters available in a virtualization environment to separate the management and production traffic of the virtual machines, these two traffic types can be lead over a common physical adapter, provided that it allows a logical separation of these two traffic types. This connects the virtual machines to only one physical layer-2 network that must not be identical to the management network of the hardware infrastructure.

Instead of direct communication, gateways between these network areas must be set up via appropriate security systems (e.g., firewalls and jumphosts).

It is recommended to use clearly distinguishable IP address ranges for these different kinds of network areas in the "1st level segregation".

Motivation: This way, systems are separated according to their purpose.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Unauthorized access or tapping of data
- Disruption of availability
- Unnoticeable feasible attacks

For this requirement the following warranty objectives are relevant:

ID: 3.58-10/2.4

Req 11	A network separation on the client side of systems must be implemented with the same separating effect on the management network as well as all other interfaces.

Systems and system components must be separated with an identical effect at all points via which communication is possible. A network-based separation of systems (e.g., through VLANs, private VLANs or other layer-2 techniques) must take place in all connected network areas (e.g., production network, management network, storage/backup network) according to the same logic. The same applies to other network technologies (e.g., Fiber Channel (FC)).

Special attention must be paid to system management and backup systems: In many cases, multiple systems are managed via a single instance. Please note that it must not be possible to communicate between systems or have one system compromise another one using the management network. So externally reachable systems must not be managed together with purely internal machines within the same network segment.

Motivation: A separation with the intention to prevent other systems from being attacked following a compromise only generates genuine added value if it is executed on all interfaces in the same way.

Implementation example: From a "customer perspective", presentation and database layer are separated on the network side. However, the administrative management interfaces would be connected in the same network and communication among systems would be possible. If the presentation layer was compromised, an attacker could access the systems of the database layer on layer-2 by circumventing a possible firewall, which regulates the communication to the data / production network.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Disruption of availability
- Unnoticeable feasible attacks

For this requirement the following warranty objectives are relevant:

ID: 3.58-11/2.4

Req 12	“Site-to-Site” connections, e.g., generic IPSec tunnels connecting internal networks between data centers, must be terminated in such a way that the tunneled communication can be checked at the relevant firewall / the packet filter of the target network.

This also concerns generically configured tunnels/VPNs between data centers which are used by multiple systems. For example, untrustworthy network sections are bridged this way.

Motivation: Encrypted connections on the IP level must be terminated in such a way that it is still possible to check the communications matrix.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Unauthorized use of services or resources

For this requirement the following warranty objectives are relevant:

ID: 3.58-12/2.4

Req 13	At the network boundary of the cloud, the datacenter or other connected infrastructure, IP packets with logically invalid addresses must be dropped both incoming and outgoing.

Packets with spoofed source addresses and invalid packets must be dropped early in the network. This is relevant for the whole datacenter network and is independent from the protection of individual systems by specialized filter elements.

Motivation: Protection of systems in the datacenter against wrong requests and malicious IP packets. Protection of foreign systems against invalid packets sent out, e.g. to avoid denial of service attacks with forged source IP addresses.

Implementation example: Denying packets using private IP addresses (as source or destination) at the border gateway to the public internet; dropping packets with source addresses which are locally to the datacenter network.

For this requirement the following threats are relevant:
- Disruption of availability

For this requirement the following warranty objectives are relevant:

ID: 3.58-13/2.4

5. Firewalls

Req 14	Filter elements like firewalls or loadbalancer must be independent from the systems to be protected.

Filter elements may be, in addition to traditional firewalls (physical or virtual), load balancers, session border controllers, access lists on routers or mechanisms implemented in a virtualization environment. This protection must not be changed directly by the system to be protected.

?For virtualized environments, it is recommended to run these firewalls on different host systems, and not on hosts for the systems to be protected. When using filtering mechanisms directly implemented in the virtualization layer, e.g., micro segmentation, the filtering takes place directly in the hypervisor of the corresponding virtual machine.

Sometimes, loadbalancer take on a double role, when used for application balancing and also acting as a filter on IP level.

Motivation: The separation of the security elements

- *ensures a clear separation between security functions and services,*
- *supports a distinction between administrative tasks and responsibilities, and*
- *prevents an attacker on a compromised system from disabling the (network-based) filters.*

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Unauthorized use of services or resources

For this requirement the following warranty objectives are relevant:

ID: 3.58-14/2.4

Req 15	Filter elements like firewalls must not contain services which have negative influence to the reliability of the firewall functionality.

Network functions which can affect each other must be implemented on different systems:
- Router using dynamic routing protocols,
- Firewalls if changes of the ruleset often occur or in case of deep packet inspection.
However, static IPSec tunnels can be terminated on a firewall if all access rules still applies and no rules are circumvented.

Motivation: More features in a single system lead to more risks of a disruption of the main functions.

For this requirement the following threats are relevant:
- Disruption of availability

For this requirement the following warranty objectives are relevant:

ID: 3.58-15/2.4

Req 16	The default configuration of a filter element (firewall) must be "deny any any", that is every needed communication has explicitly to be allowed.

Work must be based on whitelists and must be configured in such a way that, where the set of rules is processed sequentially, there is always a general deny at the end (default deny). Incoming as well as outgoing traffic must be filtered where in individual cases this may be omitted. Logically incorrect IP packets, e.g., with clearly forged sender addresses (“spoofing”) must be discarded. The configured rules must correspond to a documented communication

matrix.

To avoid activations for huge TCP/UDP port ranges, necessary services/protocols with dynamic port assignment should be investigated for alternatives, as far as a solution by using an appropriate firewall system with corresponding protocol support is not available. Otherwise, services might be accessible via the enabled ports which should actually be blocked.

Not every source IP address must be individually configured if the protection of accessible destination services is the main goal. With relation to the criticality of the systems and their need of protection IP ranges instead of single IP addresses can be configured if the maintenance overhead and the number of possible configuration errors can be significantly reduced and no higher risk to the destination systems arises.

Motivation: Despite taking great care, the use of blacklists may cause communication links to be released unintentionally. This measure is used to directly protect the systems and also to protect the entire network and ensure its stability.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Unauthorized use of services or resources
- Unnoticeable feasible attacks
- Attacks motivated and facilitated by information disclosure or visible security weaknesses

For this requirement the following warranty objectives are relevant:

ID: 3.58-16/2.4

Req 17	Filter elements (firewalls) may forward to internal systems only packets from allowed internal networks.

IP packets with addresses (source or destination) not used in internal networks have to be dropped.

Motivation: Protection against malicious packets which could influence the availability and integrity of the destination system.

Implementation example: No default route on internal systems (network elements just as servers). This ensures that only packets matching explicit existing routing entries are delivered.

For this requirement the following threats are relevant:
- Disruption of availability

For this requirement the following warranty objectives are relevant:

ID: 3.58-17/2.4

Req 18	If IPv6 is in use, filtering rules for IPv4 and IPv6 must be configured similarly.

Packet filters, firewall rules etc. must be set up so that the same protection measures apply to IPv4 and IPv6. Services that are not accessible via IPv4 must also not be accessible via IPv6.

Motivation: The use of IPv6 must not circumvent protection measures implemented for IPv4.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Unauthorized use of services or resources
- Disruption of availability

For this requirement the following warranty objectives are relevant:

ID: 3.58-18/2.4

6. System management

Req 19	The management of the datacenter or cloud infrastructure must take place via an own management network which is strictly separated from other networks.

The access for administrative tasks must be configured in such a way that only administrators are able to access the system. The access must be configured so that it is only possible via dedicated jumphosts.

It is recommended that all manageable ("active") components are always reachable within their management network. Servers via ILO, network elements like router and switche over a dedicated management port. To ensure that access is possible all the time, the management network should be part of an own network infrastructure independent from data and customer traffic. Furthermore, it is recommended to use a dedicated management network (e.g. separate VLAN) for every kind of technology and this network is reachable only from the designated management systems.

The system management of infrastructure components that provide both internal and external services must be implemented via a dedicated or the internal management network.

Motivation: Restricting possible access to a unique defined path and number of options makes it possible to clearly control access. Implemented protection mechanisms of the systems cannot be circumvented this way. The operator must not loose control over the infrastructure.

Implementation example: One VLAN for managing all internal network components and a second one for external network components. Another VLAN for the connection of the ILO ports of internal server and one for the storage systems. Communication between these VLANs must not be possible; all thess systems can only communicate with the associated management systems, possibly also with jumphosts if necessary.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Unnoticeable feasible attacks

For this requirement the following warranty objectives are relevant:

ID: 3.58-19/2.4

Req 20	Systems for system management must be used in a controlled manner.

Management systems, backup as well as monitoring systems are to be planned, installed and operated carefully. Usually they have privileged access to many systems by default. It if is not possible to do without this privileged access (e.g., with administrator rights), measures must be implemented to minimize any involved risks regarding the managed systems (“controlled”). The following applies:

- Communication between management and managed systems must be unidirectional with only one TCP/IP port.
- The initiation of communication to/from used agents depends on the trust relation of the involved systems and the transmitted data.
- Central systems should be able to do without any privileged access to the managed systems, if possible.
- A sufficient authentication of agents to the management server and vice versa is required.
- No central component should have privileged access to a critical mass of systems.
- Automated installation and configuration mechanisms should be secured in such a way that they cannot be misused for unwanted changes to systems; e.g., they should be deactivated after installation and initial configuration. Especially for installation of software packages and configuration changes, a signature mechanism could be used to ensure that only authorized changes are applied.
- Backup data must be accessible from authorized systems only.

Motivation: If a central system is compromised, the possibility to directly compromise further systems with privileged rights must be restricted.

Implementation example: To configure a system or to collect configuration parameters, a way has been chosen where the target system establishes the communication. Access in near-real-time can take place by regularly triggering. HTTPS is recommended because only one port is needed and confidentiality and authentication is included.

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Unauthorized access or tapping of data
- Unauthorized modification of data
- Disruption of availability

For this requirement the following warranty objectives are relevant:

ID: 3.58-20/2.4

Req 21	System management software for managing different types of infrastructure elements or application systems must have a role based model that is oriented to the responsibilities of the different managing tasks.

Besides a different management for different kinds of physical systems, e.g., storage or server, the management software should also separate the management functions for server and network administration in virtual environments.

Motivation: A differentiation of network and system administration supports the principle of "separation-of-duties".

For this requirement the following threats are relevant:
- Unauthorized access to the system
- Unauthorized modification of data

For this requirement the following warranty objectives are relevant:

ID: 3.58-21/2.4