INTERNATIONAL STANDARD

ISO/IEC 21878

First edition
2018-11

Information technology — Security techniques — Security guidelines for design and implementation of virtualized servers

Technologies de l'information — Techniques de sécurité — Lignes directrices pour la conception et l’implémentation sécurisées des serveurs virtualisés

Reference number
ISO/IEC 21878:2018(E)

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Fax: +41 22 749 09 47
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents) or the IEC list of patent declarations received (see http://patents.iec.ch).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see www.iso.org/iso/foreword.html.

This document was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 27, *IT Security techniques*.

Any feedback or questions on this document should be directed to the user’s national standards body. A complete listing of these bodies can be found at www.iso.org/members.html.

Introduction

Data centre infrastructures are rapidly becoming virtualized due to increasing deployment of virtualized servers (VSs) for cloud computing services and for internal IT services. Since VSs are compute engines hosting many business-critical applications, they are key resources to be protected in virtualized data centre infrastructure. As VSs are becoming mainstream in typical data centre infrastructure setups, the secure design and implementation of VSs forms an important element in the overall security strategy.

The purpose of this document is to provide security guidelines for the design and implementation of VSs. The motivation for this document is the global trend in enterprises and government agencies deploying server virtualization technologies within their internal IT infrastructure as well as the use of VSs by cloud service providers. Hence the target audience is any organization using and/or providing VSs.

The intended goal of this document is to facilitate informed decisions with respect to architecting VS configurations. Such design and implementation configuration is expected to assure the appropriate protection for all virtual machines (VMs) and the application workloads running in them in the entire virtualized infrastructure of the organization.

Information technology — Security techniques — Security guidelines for design and implementation of virtualized servers

1 Scope

This document specifies security guidelines for the design and implementation of VSs. Design considerations focusing on identifying and mitigating risks, and implementation recommendations with respect to typical VSs are covered in this document.

This document is not applicable to: (see also 5.3.2 Exclusions)

— desktop, OS, network, and storage virtualization; and

— vendor attestation.

This document is intended to benefit any organization using and/or providing VSs.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 17788, *Information technology — Cloud computing — Overview and vocabulary*

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC 17788 and the following apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

— ISO Online browsing platform: available at https://www.iso.org/obp

— IEC Electropedia: available at http://www.electropedia.org

3.1
domain
information domain
collection or cluster of *virtual machines* (3.7) hosted in one or more VSs

3.2
guest operating system
guest OS
operating system that runs within a *virtual machine* (3.7)

3.3
host operating system
host OS
operating system onto which virtualization software (hypervisor) is installed

Note 1 to entry: Host OS is an optional component of a virtualized server.

3.4
hypervisor
computer software that creates and runs one or more *virtual machines* (3.7)

3.5
management subsystem
component of a VS that enables administrators to configure the VS components

3.6
hardware
physical resources including processors, memory, devices, and associated firmware

3.7
virtual machine
VM
software-defined complete execution stack consisting of virtualized *hardware* (3.6), operating system (guest OS), and applications

3.8
virtualized server
VS
physical host on which a *hypervisor* (3.4) is installed to enable execution of multiple *virtual machines* (3.7)

3.9
virtualized server administrator
VS administrator
person with rights and responsibilities to configure and manage VS components

3.10
virtual machine manager
VMM
all software that enables VMs to run on a virtualized platform which includes the hypervisor, service VMs, virtual and physical device drivers

3.11
golden image
golden VM image
master copy of disk image of a VM or server with a specific configuration

4 Symbols and abbreviated terms

API	Application programming interface
CLI	Command line interface
COBIT 5	Control objectives for information and related technologies
CSC	Cloud service customer
CSP	Cloud service provider
IT	Information technology
NFV	Network function virtualization
OASIS	Organization for the advancement of structured information systems
OS	Operating system

PII	Personally identifiable information
SDN	Software-defined networking,
SSD	Solid state drive
VLAN	Virtual local area network

5 Overview of server virtualization

5.1 Types of server virtualization

Server virtualization is the abstraction of the underlying hardware resources to enable multiple computing stacks (consisting of an OS, middleware and applications) to be run on a single physical host. A VS is commonly classified along two perspectives.

The first perspective addresses full virtualization versus para-virtualization.

- **Full virtualization:** The guest OS is unaware that it is in a virtualized platform. The hypervisor exposes the interface of a hardware device that is physically available to the VM and for which drivers are available for guest OS, and it completely emulates the behaviour of that device. Emulation allows the programs running in VMs to use the VM OS drivers that were designed to interact with the emulated device without installing any special driver or tool specified by the hypervisor vendor.

- **Para-virtualization**: In this implementation, the hypervisor exposes a device that does not physically exist, which is just software only, and presents a lightweight interface. However, this scenario calls for having special drivers in the VM, requiring modification to the guest OS which becomes para-virtualized-aware. This approach is intended to increase the performance level of the applications running in the VM, compared to the full emulation approach adopted in full virtualization.

The second perspective is based on the platform on which the hypervisor is installed. There are two types of hypervisors:

- **Type 1** hypervisors, also known as bare-metal hypervisors, are installed directly over computer hardware, with no need for an underlying host OS.

- **Type 2** hypervisors are hosted on top of a host OS. Type 2 hypervisors are started like a regular software application before any VMs can be run and controlled.

5.2 Components of a VS

Server virtualization in the context of this document relates to a VS that implements virtualized hardware components on server-class hardware. It creates a virtualized hardware environment (virtual machines or VMs) for each instance of a guest OS permitting these environments to execute concurrently while maintaining the appearance of isolation and exclusive control over assigned computing resources. Each VM instance supports applications such as file servers, web servers, and mail servers. Server virtualization can also support client OSs in a virtual desktop or thin-client environment.

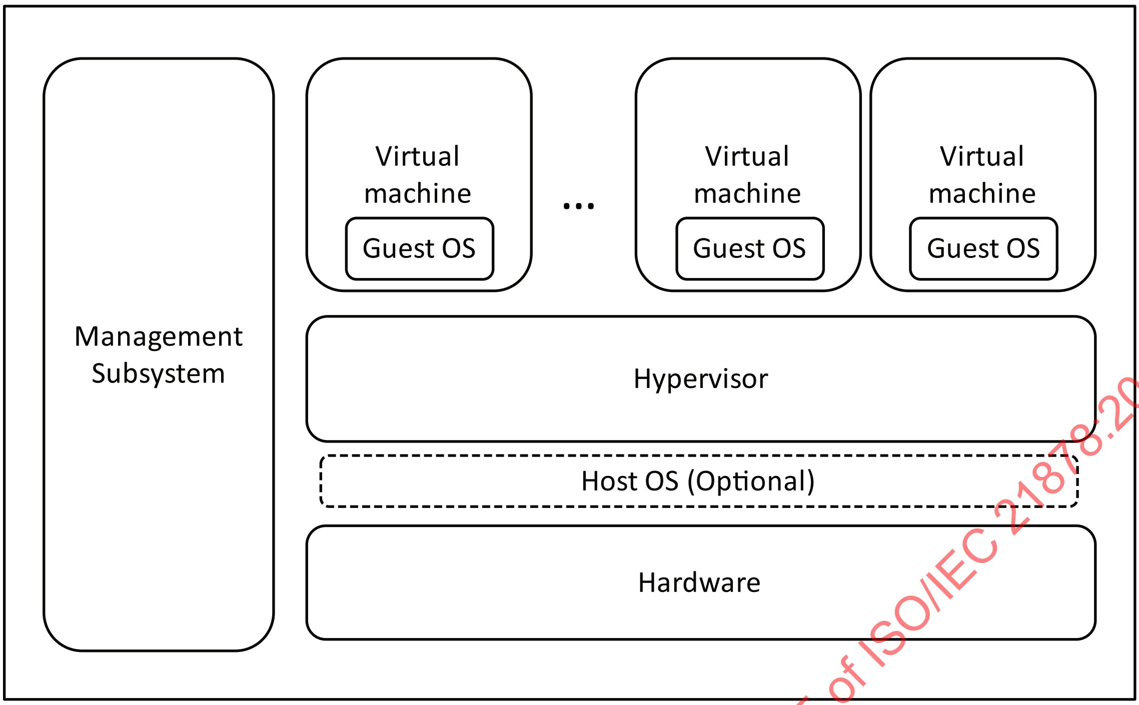

Figure 1 — Functional components of a VS

A VS comprises the following functional components as illustrated in Figure 1:

- a hypervisor;
- virtual machine(s);
- hardware;
- host OS (optional);
- other components such as management subsystem(s) for the VS administrator to configure and manage the VS.

In general, for server virtualization products that are installed onto "bare metal," the entire set of installed components constitutes the VS, and the hardware constitutes the platform. Also for products that are hosted by or integrated into a commercial off-the-shelf (COTS) OS, the components installed expressly for implementing and supporting virtualization are in the VS, and the platform comprises of the hardware and host OS.

5.3 Technical considerations

5.3.1 General

This document's recommendations are applicable to both perspectives of server virtualization as described in 5.1. From this viewpoint, the security guidelines in this document can be replicated across a larger scale (for example, in environments with thousands of VS and more).

5.3.2 Exclusions

5.3.2.1 Vendor attestation

A manufacturer or vendor of a product can claim an attestation of conformance to a standard.

This document provides guidance regarding the design and implementation of existing products and makes no specific requirements on the virtualization or associated products mentioned.

5.3.2.2 Operating environment

There are specific conditions that are assumed to exist in the operational environment where VS is deployed and hence not addressed in this document. These assumptions include both practical realities in the development of the VS security requirements and the essential environmental conditions on the use of the VS.

- **Physical security:** Physical security should be commensurate with the value of the VS and the data it contains.
- **Platform integrity**: The platform has not been compromised prior to installation of the VS.
- **Trusted administrators:** VS administrators are trusted to follow and apply all administrator guidance.

6 Overview of security threats and risks

6.1 General

The threats and risks in a virtualized infrastructure (where the computing nodes are predominantly VSs) can be classified into two types:

- Common threats: Threats common to all types of IT infrastructures (virtualized or non-virtualized);
- VS-specific risks: Risks that impact the confidentiality, integrity and availability of VSs.

6.2 Common threats

The following threats are commonly faced in virtualization setups in data centre infrastructures or cloud computing environments, but not specific to just for VSs technology per se.

- **Administrator error:** An administrator can unintentionally install or configure the VS incorrectly, resulting in ineffective security mechanisms.
- **Data leakage:** If it is possible for data to leak between domains when prohibited by policy, then an adversary on one domain or network can obtain data from another domain. Such cross-domain data leakage can, for example, cause classified information, corporate proprietary information, or PII data to be made accessible to unauthorized entities.
- **Insecure network configuration:** The VS is itself a node of a larger enterprise network and hence insecure network location and insecure host-level protection can have an impact on all the software and applications running in it.
- **Platform compromise:** The hosting of untrusted or malicious domains by the VS can compromise the security and integrity of the platform on which the VS executes.
- **Poor cryptography key management:** When the keys of the cryptographic applications are stored in insecure locations, the security of encrypted data can easily be compromised.
- **Third-party software:** Vulnerabilities in third-party software used as a VS component (e.g., device drivers) can compromise the security of an entire VS.
- **Unauthorized access:** A user can gain unauthorized access to the VS data and VS executable code. A malicious user, process, or external IT entity can masquerade as an authorized entity in order to gain unauthorized access to VS data or VS resources.

- **Unauthorized modification:** Malware running on the physical host can undetectably modify VS components while the components are in execution mode or at rest. Likewise, malicious code running within a VM can modify VS components.
- **Unauthorized update:** A malicious program can gain administrative privileges to perform an unauthorized update that can compromise the integrity of one or more VS functions.
- **VMM compromise:** Failure of security mechanisms can lead to unauthorized intrusion into or modification of the VMM or bypass of the VMM altogether.
- **Weak cryptography:** A threat of weak cryptography can arise if the VMM does not provide sufficient entropy to support security-related features that depend on entropy to implement cryptographic algorithms.

6.3 VS-specific risks

6.3.1 General

VS-specific risks are examined in further detail here. Moving to a virtualized environment brings a unique set of risks not present in traditional IT environment. These risks include:

- VM risks;
- hypervisor risks;
- operational risks related to implementation;
- cloud services risks.

6.3.2 VM risks

The use of VMs can either introduce new and unique security risks or lead to more significant impacts for particular known risks. Consequently, as part of assessing the risks of virtualization, the following should be considered:

- **VM sprawl**: Uncontrolled proliferation of VMs can lead to an unmanageable condition of unpatched and unaccounted for machines. In a traditional IT environment, physical servers are procured. This requirement enforces effective controls, because change requests should be created and approved before hardware and software can be acquired and connected to the data centre. In the case of virtualized environments, however, VMs can be allocated quickly, self-provisioned, or moved between physical servers, bypassing the conventional change management process. Without an effective control process in place, VMs and other virtual systems with unknown configurations can quickly proliferate, consuming resources, degrading overall system performance, and increasing liability and risk of exposure. Because these machines are typically not readily detectable or visible, they cannot be effectively monitored or tracked for the application of security patches or effectively investigated when a security incident occur.
- **Sensitive data within a VM:** Data confidentiality within VMs can be easily compromised, because data can easily be transported and tampered with. Although VM images and snapshots provide a way to deploy or restore virtual systems quickly and efficiently across multiple physical servers, this capability means that copies of images and snapshots can be removed from a data centre easily. This removal includes current memory contents, which are not intended to be stored on the storage devices themselves. Therefore, in a virtualized environment, it is no longer possible to assure that sensitive data such as system password files is safe from unauthorized personnel. This sensitive information and the VM containing it can be moved easily, making it possible to compromise a VM and reintroduce it into the system later. Without proper controls, security loopholes can exist through the inadvertent capture, storage, and deployment of sensitive information, including rogue virtual instances of critical services. Potential hackers or disgruntled staff can gain access and insert malicious code into VM images and snapshots, which can then be rapidly deployed throughout the environment, resulting in its compromise.

- **Security of offline and dormant VMs:** Dormant and offline VMs can eventually deviate so far from a current security baseline that simply powering them on introduces massive security vulnerabilities. Dormant or offline VMs can easily be overlooked and left out of the execution of important security procedures. For example, it is likely that a dormant VM is not updated with the latest security patches. As a result, when the VM is run again, it can be exposed to the latest vulnerabilities. Similarly, dormant VMs can also lack up-to-date access control policy or be excluded from essential security monitoring functions, which can cause security loopholes in the virtualized environment. Critical infrastructure, management, and security services are increasingly packaged and delivered as virtual appliances. These appliances require proper classification and treatment to prevent rogue instances of non-compliant policies and configurations. In addition, if these appliances are not limited to specific clusters of physical hosts or storage, exposure to risk can increase.

- **Security of pre-configured (golden image) VM/active VMs:** Virtual machines existing as files on a virtualization platform can easily be transported via physical means or through a network. This can lead to unauthorized access, resulting in machine configuration changes or viral payload injection into the platform's virtual disks. Unauthorized access through malicious interception can compromise these VM images. Often, golden VM images are made, making it easy to deploy cloned copies and to compromise the integrity of the virtualized environment.

- **Lack of visibility into and control over virtual networks:** Software-defined virtual networks can cause network security breaches, because traffic over virtual networks cannot be visible to security protection devices on the physical network. Traditional IT infrastructure network traffic is monitored and secured as data flows across actual routers, switches, and firewalls. On a virtual network, this can become unmanageable unless the traffic is explicitly redirected to physical or virtual appliances for monitoring. Virtual network configuration can be modified with relative ease; this capability can cause conflicts with actual physical network security policies.

- **Resource exhaustion:** Uncontrolled physical resource consumption by virtual processes can lead to reduced availability. Resource-intensive software tends to exhaust resources in a physical server when it is implemented in multiple VMs. For example, anti-virus and other security software interrupt every call to disk or memory in order to monitor and prevent security incidents such as hacking or viruses. When anti-virus software runs simultaneously in different VMs on the same physical server, it can potentially consume the host resource pool. Similarly, automated OS patches on a large group of VMs can have the same effect.

6.3.3 Hypervisor risks

A risk factor unique to virtual environments is the hypervisor. The hypervisor is the software and/or firmware responsible for hosting and managing VMs. It provides a single point of access into the virtual environment and is also potentially a single point of failure. A misconfigured hypervisor can result in a single point of compromise of the security of all its hosted components. It does not matter how individual VMs are hardened — a compromised hypervisor can override those controls and provide a convenient single point of unauthorized access to all the VMs. The following security risks are related to the use of hypervisor:

- **Securing the hypervisor:** It is the process of ensuring that the hypervisor, the software that enables virtualization, is secure throughout its life cycle, including development, implementation, provisioning, and management.

  - A hypervisor can control all aspects of VM operation, so it is a natural target of malicious attacks. Securing a hypervisor is vital, yet more complex than it seems. In an attack known as "hyper-jacking," malware that has penetrated one VM can attack the hypervisor. When a guest VM attempts this attack, it is often called a "guest VM escape," because the guest VM breaks out of its isolated environment and attacks the host hypervisor. Once compromised, the hypervisor can be used as an attack platform to compromise guest VMs hosted by it or other hypervisors with which it can be able to interact.

  - All hypervisors in common use include a very rich set of remote management API and hence their attack surfaces are increased. Calling tools/scripts/applications not adequately implementing identity and access control, especially if a hypervisor uses locally managed service accounts, present greater risks.

- **Unauthorized access to hypervisor** Administrative access controls to the hypervisor can be insufficient to protect against potential hacker attacks.
  - The hypervisor can be accessed using a local (host level) authentication scheme or enterprise level authentication scheme. Local authentication schemes especially those using passwords can provide only weak authentication and cannot support robust enterprise security policies and associated authentication schemes.
  - Hypervisors in a data centre can be managed individually or centrally through an enterprise virtualization management software. In the latter approach, if the management software communicates with hypervisor without adequate integrity and network level protections, attackers can exploit this situation to gain unauthorized administrative control to multiple hypervisor hosts.

6.3.4 Operational risks related to implementation

Compared to traditional IT environments, virtualization of IT systems inevitably leads to changes in operational procedures. As a result, some common defence-in-depth practices used in securing physical servers can be affected or ignored, while newly introduced features or functions can expose the environment to additional risks. The following are the security risks related to changes in operation procedures.

- **Account or service hijacking through the self-service portal:** Portal vulnerabilities can lead to privilege escalation attacks. A self-service portal is typically used to delegate specific parts of virtual infrastructure provisioning and management to assigned self-service administrators. Liberal use of self-service portals in cloud computing services increases susceptibility to security risks, including account or service hijacking.
- **Workloads of different trust levels located on the same server:** Ensure that there is sufficient security segregation of workloads on a physical host. Enterprises can attempt to segregate VMs of different trust levels on separate host machines. However, that effort is effective only if there is effective implementation of systems and data categorization, as well as implementation of adequate network, security, and management controls. VMs of lower trust levels typically have weaker security controls than VMs of higher trust levels. Those VMs can therefore be easier to compromise, potentially providing a stepping-stone to higher-risk, more sensitive VMs on the same host. It is important to have consistency in the levels of protection for VMs of different trust levels across physical and virtual environments. In short, hosting VMs of different trust levels on the same host tends to reduce overall security for all components to that of the least-protected component.

6.3.5 Cloud Services risks

Enterprise IT personnel can interact with virtualization technologies when they are using cloud services from CSPs. In such cases, it can introduce additional risk factors, including the following:

- **Risk due to cloud service provider APIs:** A hybrid (private/public) cloud virtualization implementation can pose security risks due to account/authentication federation. CSPs expose a set of software interfaces or APIs that an enterprise can use to manage and interact with cloud services. Such interfaces when not designed properly to protect against accidental and malicious attempts to circumvent enterprise policies, can pose additional risks.
- **Use of VMs within cloud services:** Many CSPs offer services that are based on VM functionality. As outlined in ISO/IEC 19941, the security requirements and the responsibilities for implementing these requirements vary depending on the cloud service capabilities type. For example, in an infrastructure cloud capabilities type of cloud service, the CSC can be responsible for configuring and implementing various aspects of VM security (and in some instances, hypervisor security), whereas some of this responsibility can fall to the CSP in the case of a platform cloud capabilities type of cloud service.

- **General cloud services security:** When securing VMs in a cloud environment, other aspects of security should also be considered. Further information on this topic can be found in ISO/IEC 19086-4, ISO/IEC 19941, ISO/IEC 27017 and ISO/IEC 27018.

7 Recommendations for secure VS lifecycle

7.1 General

The objectives of secure VS design and implementation are to assure the confidentiality, integrity and availability of information processed and stored in the virtual hosts. Secure deployment of VSs requires several strategic tasks to be carried out in all the four typical phases of deployment, i.e., initial preparation, planning and design, implementation and disposition. These tasks are described in 7.2 to 7.5. In addition, Clause 8 provides security issues for consideration, and Clause 9 provides a checklist for the implementation of security of VSs.

When an organization embarks on a server virtualization initiative, it should ensure that its information security governance framework also applies to its virtualized IT systems and services. Delivering enterprise stakeholder value through virtualization initiatives requires good governance and management of IT assets. Organizations that choose to virtualize should opt for a comprehensive framework, such as COBIT 5, that enables them to meet their technology goals and deliver value.

An integral part of security governance framework is the risk assessment process. The risk assessment process in the context of a virtualized infrastructure consists of identifying the risks (see Clause 6), estimating the likelihood of their occurrence and assessing their impact in order to establish appropriate controls to address them before implementation. ISO/IEC 27001 and ISO/IEC 27005 provide more details on a process that can be used or adapted by enterprises of various sizes and complexities. Some of the key elements to be considered when performing a risk assessment for virtualized infrastructure can be found in Annex A.

An organization should establish policies and procedures that include an audit program geared to virtual IT systems. Roles and responsibilities of system administrators and users should be clearly defined and documented. An organization should govern a virtualization initiative by evaluating, directing, and monitoring every step in the process. In this context, IT managers should assure that their teams follow virtualization policies and procedures holistically across the enterprise.

7.2 Initial preparation phase

During the initial phase, an organization should identify virtualization needs, providing an overall vision for how virtualization solutions support its mission; creating a high-level strategy for implementing virtualization solutions; developing virtualization policy and identifying platforms and applications that can be virtualized; and specifying business and functional requirements. The preparation work for the design and the implementation of VS security involves the following stages:

- asset identification;
- requirements collection;
- review of requirements;
- evaluation of technical options and constraints;
- evaluation of existing designs and implementations with respect to security requirements.

ISO/IEC 27033-2 provides more details on how these five stages provide inputs to the design and implementation of network security; however, these stages are also applicable to the design, implementation and disposition of VS security.

7.3 Planning and design phase

During the planning and design phase, an organization should provide necessary guidance for specifying and evaluating the technical characteristics of the virtualization solution and related components, including authentication methods and cryptographic mechanisms to protect communications. Major considerations include selection of virtualization software including the security functionality of the hypervisor and its assurance level, storage system, network topology, bandwidth availability and deployment environment (cloud service or on premises). In particular, the assurance is required for addressing all categories of risk in 6.3.3 and 6.3.5 as applicable to the deployment environment. As an example, the set of certified hardware platforms on which the targeted hypervisor can be run and the hardware support for virtualization by those platforms are key factors in reducing the runtime vulnerabilities of the hypervisor. The design should also take into account the appropriate logical segregation of instances that contain sensitive data.

The management consoles or their equivalent that are part of hypervisor design should support secure communication protocols. Separate authentication should be established for application/server, guest OS, hypervisor, and host OS to provide different layers of security and protection. Access control policies for VS and VMs should be formulated. An organization should also define and document processes for handling incidents that involve virtualization solutions.

Security considerations for planning and design phase are described in 8.2.

7.4 Implementation phase

During implementation, an organization should assure that sound security practices are established through extensive assessment of the vulnerability of the virtualization components. The underlying virtualization platform should be hardened using vendor-provided guidelines and/or third-party tools.

In a virtualized environment, robust key management is essential to access control and proof of ownership for both data and keys. Role-based access policies should be enforced to enable segregation of duties, thereby facilitating proof of governance. Proper data governance measures are required to identify, track, and control where data instances containing sensitive assets reside at any given time. The hypervisor should have introspection capabilities to monitor VM behaviours and configuration capabilities to enforce VM operations for conformance to chosen security policies. Proper encryption of VM files is required to significantly reduce the risk associated with unauthorized user access to physical servers and storage containing sensitive data. Implementation plan should include tasks such as initial and on-going configuration review and security testing for side channel vulnerabilities so that data leakage and unauthorized disclosure of sensitive data are prevented. Automated tools for timely patching of hypervisor vulnerabilities and monitoring of critical configuration files need to be considered at this phase.

A security checklist for implementation phase is described in 9.2.

7.5 Disposition phase

The consideration for the disposition phase should begin early in the lifecycle for the secure deployment of VSs. During the disposition phase, tasks should be clearly defined as regards sanitizing media before disposition. The VM retirement process should prevent data leakage and breaches, including shredding or revocation of keys associated with encrypted VMs. Periodic internal and external audits of the virtualized environment facilitates early identification and mitigation of weaknesses and vulnerabilities.

8 Planning and design phase: security considerations

8.1 General

The security considerations that enable the recommendations to be carried out for planning and design phase (7.3) are given in Table 1. A brief description of each security consideration is also given in the table.

8.2 Security considerations and satisfying requirements

Table 1 — Security considerations (Planning and design phase)

No.	Security consideration	Description
1	Can/Is the VM be isolated?	As basic functionality, the VMM should support a security policy that mandates no information transfer between VMs. Configure VMs such that each executes in its own space and does not need to share resources and data with other VMs resident in the VS.
2	Is the integrity of the VMM in the VS taken care of?	Integrity is a core security objective for VS. To achieve system integrity, the integrity of each VMM component should be established and maintained. This objective concerns only the integrity of the VS — not the integrity of software running inside of VMs or of the physical platform. The overall objective is to assure the integrity of critical components of a VS. Ensure that VMM is not modified or bypassed so that its ability to enforce isolation of VMs is not compromised.
3	Is the integrity of the platform taken care of?	The integrity of the VMM depends on the integrity of the hardware and software on which the VMM relies. Although the VS does not have complete control over the integrity of the platform, the VS should as much as possible try to assure that no users or software hosted by the VS is capable of undermining the integrity of the platform. Ensure the integrity of the hardware (and optionally the host OS) on which the VS components are installed, since if the integrity of the platform is compromised, the entire VS can be compromised.
4	Is the access control for management functions in the VS planned?	Only authorized administrators should exercise management functions. Ensure that no unauthorized access is allowed for VS management functions.
5	Is the VS designed with manageability?	The VS should support protocols that facilitate management of the VS as an IT product. Ensure that there is API support for carrying out all necessary management functions.
6	Can the VS be audited?	The purpose of audit is to capture and protect data about what happens on a system so that it can later be evaluated at a later time. Ensure that logs of all events and actions on VS and resident VMs can be created in a form that facilitates analysis both in real time and post facto.
7	Are the administrative configurations set according to the security policy?	VS administrators should configure the VS and ensure that there are flexible configuration options available to configure the VS for a wide range of policy requirements.
8	Is physical security ensured?	The environment provides physical security, commensurate with the value of the VS and the data it contains. Ensure that there are adequate physical security safeguards both for normal and contingency operations.
9	Will techniques increasing trust of administrators be implemented?	Capability to ensure that there are processes for recruitment and training that guarantees trustworthiness of administrators.

9 Implementation phase: security checklist

9.1 General

The security checklist that addresses each of the potential risks identified in Clause 6 and the vulnerability exposure due to not implementing the associated measures is given in Table 2. Collectively all the items in the checklist provide coverage for the strategic tasks to be carried out for implementation phase (7.4). The implementation guidelines for implementing these security checklist items are given in Annex B.

9.2 Security checklist and vulnerability exposure

Table 2 — Security checklist (Implementation phase)

No.	Security checklist item	Addressed risk and vulnerability exposure due to non-implementation
1	Have the policies and controls which prevent uncontrolled proliferation of VMs been implemented?	Risk addressed: VM sprawl (6.3.2) Vulnerability exposure: An uncontrolled proliferation of VMs can lead to an unmanageable condition of unpatched and unaccounted machines.
2	Are the sensitive data within the VM being protected?	Risk addressed: Sensitive data within a VM (6.3.2) Vulnerability exposure: Failure to consider the protection of sensitive data within the VS can lead to data easily transported and tampered with. The organization can seek a solution that incorporates effective, policy-based key management of data stored in physical, virtual, and cloud servers.
3	Are the security measures implemented for offline and dormant VMs?	Risk addressed: Security of offline and dormant VMs (6.3.2) Vulnerability exposure: Dormant and offline VMs can eventually deviate from a current security baseline that simply switching them on introduces potential security vulnerabilities.
4	Are the security measures implemented for pre-configured (Golden Image) VM and active VMs?	Risk addressed: Security of pre-configured (golden image) VM/active VMs (6.3.2) Vulnerability exposure: VMs exist as files on a virtualization platform and can be easily transported via physical means or through a network. However, this can lead to unauthorized access, resulting in machine configuration changes or viral payload injection into the platform's virtual disks. Unauthorized access through malicious interception can compromise these VM images. Often, golden VM images are created and stored in the image repository which are then cloned to be instantiated into VM instances. These golden VM images when not protected against unauthorized modification can result in compromised VM instances.
5	Is the visibility over traffic and controls in virtual networks ensured?	Risk addressed: Lack of visibility into and control over virtual networks (6.3.2) Vulnerability exposure: On a virtual network, visibility of network traffic can become unmanageable unless the traffic is explicitly redirected to physical or virtual appliances for monitoring. Virtual network configuration can be modified with relative ease and can cause conflicts with actual physical network security policies.
6	Are controls and policies to prevent resource exhaustion implemented?	Risk addressed: Resource exhaustion (6.3.2) Vulnerability exposure: Uncontrolled physical resource consumption by virtual processes can lead to reduced availability. Resource-intensive software tends to exhaust resources in a physical server when it is implemented in multiple VMs. For example, anti-virus and other security software interrupt every call to disk or memory in order to monitor and prevent security incidents such as hacking or viruses. When anti-virus software runs simultaneously in different VMs on the same physical server, it can potentially consume the host resource pool. Automated OS patches on a large group of VMs can have the same effect.
7	Are measures taken to ensure the security of the hypervisor?	Risk addressed: Hypervisor risks (6.3.3) Vulnerability exposure: It is critical that the organization assures that the hypervisor is secure throughout its life cycle, including development, implementation, provisioning, and management.
8	Are the measures to prevent unauthorized access to the hypervisor implemented?	Risk addressed: Unauthorized access to hypervisor (6.3.3) Vulnerability exposure: Inadequately managed administrative access controls to the hypervisor are unable to protect against potential hacker attacks.

Table 2 *(continued)*

No.	Security checklist item	Addressed risk and vulnerability exposure due to non-implementation
9	Are the controls and policies to prevent account or service hijacking implemented?	Risk addressed: Account or service hijacking through the self-service portal (6.3.4) Vulnerability exposure: Service or account hijacking via portal vulnerabilities can lead to privilege escalation attacks.
10	Are security measures for proper and secure segregation of workload on physical hosts implemented?	Risk addressed: Workloads of different trust levels located on the same server (6.3.4) Vulnerability exposure: Ensure that there is sufficient security segregation of workloads on a physical host.
11	Are measures taken to ensure the security of the cloud service provider API?	Risk addressed: Risk due to cloud service provider APIs (6.3.5) Vulnerability exposure: A hybrid (private/public) cloud virtualization implementation can pose security risks due to account/authentication federation. The API provided by cloud service providers should also be verified before implementing with the organization's cloud computing resources.
12	Are there automated processes in place to monitor patch releases and apply them to hypervisor software and Guest OS software modules?	Risk Addressed: Compromised Hypervisor in Hypervisor Risk (6.3.3) and Compromised VM in Operational risked related to implementation (6.3.4) Vulnerability Exposure: Exposure to known vulnerabilities that the patches were designed to protect.

Annex A
(informative)

Risk assessment for VSs

A.1 General

The objective of a risk assessment using Table A.1 is to evaluate the security risk of the existing (or newly designed) virtualised server and determine how best to mitigate those risks to meet the organization's needs in the following steps:

a) estimate the likelihood of the risk happening using Table A.2;

b) evaluate the severity of the impact to the enterprise due to confidentiality-integrity-availability based on Table A.3;

c) evaluate the overall severity impact based on Table A.4 using the highest value of impact due to confidentiality-integrity-availability compromise in the previous 3 columns;

d) define controls to eliminate or reduce the risk until the overall severity impact is within acceptable level;

e) evaluate the residual risk level after the controls have been applied;

f) repeat steps c) and e) for next highest impact due to confidentiality-integrity-availability until all risks have been eliminated/mitigated to acceptable levels.

A.2 Risk assessment matrix

Table A.1 provides a sample template for risk evaluation.

Table A.1 — Matrix for evaluation of risks

Type of risk	Likelihood (Table A.2)	Impact due to confidentiality compromise (Table A.3)	Impact due to integrity compromise (Table A.3)	Impact due to availability compromise (Table A.3)	Evaluate risk level (Table A.4)	Risk treatment control to be implemented	Evaluate residual risk level (Table A.4)
VM sprawl	Low Medium High	Low Medium High	Low Medium High	Low Medium High			
Sensitive data in VM	Low Medium High	Low Medium High	Low Medium High	Low Medium High			
Security of offline/dormant VMs	Low Medium High	Low Medium High	Low Medium High	Low Medium High			
Security of pre-configured (golden image) VM/active VMs	Low Medium High	Low Medium High	Low Medium High	Low Medium High			

Table A.1 *(continued)*

Type of risk	Likelihood (Table A.2)	Impact due to confidentiality compromise (Table A.3)	Impact due to integrity compromise (Table A.3)	Impact due to availability compromise (Table A.3)	Evaluate risk level (Table A.4)	Risk treatment control to be implemented	Evaluate residual risk level (Table A.4)
Lack of visibility and controls on virtual networks	Low Medium High	Low Medium High	Low Medium High	Low Medium High			
Resource exhaustion	Low Medium High	Low Medium High	Low Medium High	Low Medium High			
Hypervisor security	Low Medium High	Low Medium High	Low Medium High	Low Medium High			
Unauthorized access to hypervisor	Low Medium High	Low Medium High	Low Medium High	Low Medium High			
Account or service hijacking through self- service portal	Low Medium High	Low Medium High	Low Medium High	Low Medium High			
Workload of different trust levels located on the same server	Low Medium High	Low Medium High	Low Medium High	Low Medium High			
Risk due to CSP API	Low Medium High	Low Medium High	Low Medium High	Low Medium High			

A.3 Likelihood rating

Table A.2 provides the evaluation criteria for determining the likelihood of risk for related vulnerability.

Table A.2 — Likelihood rating for related vulnerability

Likelihood rating	Evaluation criteria
High	Relevant security control(s) is not in place.
Medium	Relevant security control(s) is in place but neither consistent nor effective.
Low	Relevant security control(s) is in place and effective.

A.4 Impact rating

Table A.3 defines the criteria for determining the impact severity of risk for related CIA compromise.

Table A.3 — Impact rating for related confidentiality-integrity-availability (CIA) compromise

Impact rating	Evaluation criteria
High	There is significant business impact to the enterprise.
Medium	There is tangible or intangible loss to the enterprise
Low	There is insignificant business loss due to minor inconvenience/inefficiency in business operations.

A.5 Risk matrix

Table A.4 defines the risk levels based on the inter-relationship between the likelihood rating (Table A.2) and the impact rating (Table A.3).

Table A.4 — Risk matrix showing the defined risk levels

Likelihood	Impact		
	Low	**Medium**	**High**
Low	1 (Insignificant)	2 (Minor)	3 (Medium)
Medium	2 (Minor)	3 (Medium)	4 (High)
High	3 (Medium)	4 (High)	5 (Very High)

Annex B
(informative)

Guidelines for implementing security checklist items in Table 2

B.1 Have the policies and controls which prevent uncontrolled proliferation of VMs been implemented?

The organization should consider the following during implementation.

- Implement effective policies, guidelines, and processes to govern and control VM lifecycle management, including self-service and automated scripts/DevOps tools.
- Control the creation, storage, and use of VM images with a formal change management process and tools. Approve additions only when necessary.
- Keep a small number of known-good and timely patched images of a guest OS separately and use them for fast recovery and restoration of systems to the desired baseline.
- Discover VMs, including dormant ones and the applications running on them, regularly. Discovering, classifying, and implementing appropriate security controls for each VM and its associated network connections is critical. This process includes a quarantine or rollback capability in case of a compromise.
- Use virtualization products with management solutions to examine, patch, and apply security configuration changes to VMs.

B.2 Are the sensitive data within the VM being protected?

The organization should consider the following during implementation.

- Encrypt data stored on virtual and cloud servers to make it unreadable. Seek a solution that incorporates simple, policy-based key management of data stored in physical, virtual, and cloud servers. Only release encryption/decryption keys to validated and authorized physical or virtual servers. Provide options to manage the keys on premises and/or in the cloud as a service. Leverage a policy-based key management system to determine where and when encrypted data can be accessed. In addition, apply identity and integrity checks when VMs request access to secure storage volumes. It is recommended that both boot and data volumes be encrypted.
- Develop policies to restrict storage of VM images and snapshots. If it is necessary to store images and snapshots, proper authorization, such as secondary level of approval, should be obtained and corresponding monitoring and control processes established. To reduce risk, carefully consider where to store these duplicate images or snapshots. For review or audit purposes, mitigation should include logging of activities, as well as establishing a formal image change management process that includes creation, distribution, storage, use, retirement, and destruction.
- Implement policies to assure that backup and failover systems, including temporary upgrade/patch instances, are cleaned when deleting and wiping (zero-filling) the VM images. Special care should be taken when using SSD drives to avoid "residual data".
- Consider using cryptographic checksum protection to detect unauthorized changes to VM images and snapshots.
- Identify critical data files within the VM that need a higher degree of monitoring as well as log management.

B.3 Are the security measures implemented for offline and dormant VMs?

The organization should consider the following during implementation.

- Control the backup, archiving, distribution, and restart of VMs with effective policies, guidelines, and processes such as suitably tagging the VM based on sensitivity/risk level.
- Use virtualization products with management solutions that examine, patch, and apply security configuration changes. While evaluating these products, consider the coverage provided across hypervisors and if there are exceptions in fine print.
- Create a controlled environment to apply security patches and control policies to an offline or dormant VM.
- Avoid problems such as VMs being accidentally or intentionally restarted or rogue instances being instantiated with measures such as appropriate architecture and design as well as regular monitoring of virtual appliances that provide critical infrastructure, management, and security services.

B.4 Are the security measures implemented for pre-configured (Golden Image) VM and active VMs?

The organization should consider the following during implementation.

- Assure proper hardening and protection of VM instances.
- Augment VM OSs with built-in security measures, leveraging third-party security technology, such as discovery and monitoring tools, to provide layered security controls.
- Consider implementing an integrity checksum mechanism for all VM images.
- Encrypt VM images to prevent unauthorized modification. At the same time, consider performance concerns depending on the data type and the underlying physical server capabilities.
- Implement strict controls and processes around access, creation, and deployment of VM images/instances.

B.5 Is the visibility over traffic and controls in virtual networks ensured?

The organization should consider the following during implementation.

- Monitor virtual networks and data traffic similarly to physical networks. Organization should carefully determine the tool to use for this task and should configure it with network port mirroring to preferably give a unified view of traffic across physical as well as virtual networks.
- Consider a hypervisor that can monitor each guest OS (VM introspection) as it is running, if separate tools are not installed to monitor communications between VMs.
- Implement security technologies that span physical and virtual environments with a consistent policy management and enforcement framework.
- Create consistent security policy and configuration across the physical/virtual network.
- Use VM-specific security mechanisms embedded in hypervisor APIs to provide granular monitoring of traffic crossing VM control and data planes. Leverage tools that implement emerging technologies such as (SDN, NFV or OpenFlow. These mechanisms are opaque to traditional network security controls.