



Red Hat Enterprise Linux 5.6 KVM Security Target

Version:	0.12
Status:	Released
Last Update:	2011-05-18
Classification:	Red Hat and atsec public

Trademarks

Red Hat and the Red Hat logo are trademarks or registered trademarks of Red Hat, Inc. in the United States, other countries, or both.

atsec is a trademark of atsec information security GmbH

Linux is a registered trademark of Linus Torvalds.

UNIX is a registered trademark of The Open Group in the United States and other countries.

IBM, IBM logo, bladecenter, eServer, iSeries, OS/400, , POWER3, POWER4, POWER4+, pSeries, System p, POWER5, POWER5+, System x, System z, S390, xSeries, zSeries, zArchitecture, and z/VM are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Intel, Xeon, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

This document is based in parts on the Red Hat Enterprise Linux Version 5 Security Target for CAPP, RBAC and LSPP Compliance, Copyright © 2007 by IBM Corp. and atsec information security corp.

Legal Notice

This document is provided AS IS with no express or implied warranties. Use the information in this document at your own risk.

This document may be reproduced or distributed in any form without prior permission provided the copyright notice is retained on all copies. Modified versions of this document may be freely distributed provided that they are clearly identified as such, and this copyright is included intact.

Classification Note

Public Information

Revision History

Revision	Date	Author(s)	Changes to Previous Revision
0.1	2010-07-30	Isabell Fouquet	Initial Version
0.2	2010-08-02	Isabell Fouquet	Host System and Virtualization Functionality
0.3	2010-08-09	Isabell Fouquet	Instantiation of SFRs
0.4	2010-08-13	Stephan Mueller	Instantiation of SFRs, Ordering of SFRs, addition of chapters 1, 2, 7
0.5	2010-08-20	Stephan Mueller	Inclusion of developer and evaluator comments
0.6	2010-08-26	Stephan Mueller	Inclusion of developer and evaluator comments
0.7	2010-10-17	Stephan Mueller	Fixing of typos, level of auditing set to "not specified", removal of FCS_RNG.1, removal of stateful inspection for IPv6, AES-CTR for SSH defined
0.8	2010-12-22	Stephan Mueller	Update description for accessing the host system, removal of PCI passthrough, ebtables update, clarification of SFRs
0.9	2011-01-14	Stephan Mueller	Update of SFR-TSS mapping, addition of certifier comments
0.10	2011-04-06	Stephan Mueller	Clarification of hardware
0.11	2011-04-07	Stephan Mueller	Consolidation of TOE type
0.12	2011-05-18	Stephan Mueller	Update FCS_RNG.1 definition to meet BSI requirements

Table of Contents

1	Introduction	10
1.1	Security Target Identification	10
1.2	TOE Identification	10
1.3	TOE Type	10
1.4	TOE Overview	10
1.4.1	Required Hardware and Software	11
1.4.2	Intended Method of Use	11
1.4.2.1	KVM virtual machines	11
1.4.2.2	Operating environment	12
1.4.3	Major Security Features	12
1.5	TOE Description	13
1.5.1	Introduction	13
1.5.2	TOE boundaries	13
1.5.2.1	Physical	13
1.5.2.2	Logical	14
1.5.2.3	Configurations	15
1.5.2.4	TOE Environment	16
1.5.2.5	Security Policy Model	16
2	CC Conformance Claim	18
3	Security Problem Definition	19
3.1	Threat Environment	19
3.1.1	Assets	19
3.1.2	Threat Agents	19
3.1.3	Threats countered by the TOE	19
3.2	Assumptions	20
3.2.1	Environment of use of the TOE	20
3.2.1.1	Physical	20
3.2.1.2	Personnel	21
3.2.1.3	Procedural	21
3.2.1.4	Connectivity	21
3.3	Organizational Security Policies	21
4	Security Objectives	23
4.1	Objectives for the TOE	23
4.2	Objectives for the Operational Environment	24
4.3	Security Objectives Rationale	25
4.3.1	Coverage	25
4.3.2	Sufficiency	26
5	Extended Components Definition	32
5.1	Class FCS: Cryptographic support	32
5.1.1	Random number generator (RNG)	32
5.1.1.1	FCS_RNG.1 - Random number generation	32
5.2	Class FDP: User data protection	33

5.2.1	Residual information protection (RIP)	33
5.2.1.1	FDP_RIP.3 - Full residual information protection of resources	33
5.3	Class FIA: Identification and authentication	33
5.3.1	User-subject binding (USB)	34
5.3.1.1	FIA_USB.2 - Enhanced user-subject binding	34
6	Security Requirements	35
6.1	TOE Security Functional Requirements	35
6.1.1	Host system related functionality	39
6.1.1.1	Audit data generation (FAU_GEN.1)	39
6.1.1.2	User identity association (FAU_GEN.2)	39
6.1.1.3	Audit review (FAU_SAR.1)	40
6.1.1.4	Restricted audit review (FAU_SAR.2)	40
6.1.1.5	Selective audit (FAU_SEL.1)	40
6.1.1.6	Protected audit trail storage (FAU_STG.1)	40
6.1.1.7	Action in case of possible audit data loss (FAU_STG.3)	41
6.1.1.8	Prevention of audit data loss (FAU_STG.4)	41
6.1.1.9	Cryptographic key generation (FCS_CKM.1(SYM))	41
6.1.1.10	Cryptographic key generation (FCS_CKM.1(RSA))	42
6.1.1.11	Cryptographic key generation (FCS_CKM.1(DSA))	42
6.1.1.12	Cryptographic key distribution (FCS_CKM.2(NET))	42
6.1.1.13	Cryptographic key destruction (FCS_CKM.4)	42
6.1.1.14	Cryptographic operation (FCS_COP.1(NET))	43
6.1.1.15	Random number generation (Class DRG.2) (FCS_RNG.1(SSH-DFLT))	43
6.1.1.16	Random number generation (Class DRG.2) (FCS_RNG.1(SSH-FIPS))	44
6.1.1.17	Subset access control (FDP_ACC.1(PSO))	44
6.1.1.18	Security attribute based access control (FDP_ACF.1(PSO))	44
6.1.1.19	Complete information flow control (FDP_IFC.2(NI))	46
6.1.1.20	Simple security attributes (FDP_IFF.1(NI-IPTables))	46
6.1.1.21	Simple security attributes (FDP_IFF.1(NI-ebtables))	47
6.1.1.22	Import of user data with security attributes (FDP_ITC.2(BA))	48
6.1.1.23	Full residual information protection (FDP_RIP.2)	49
6.1.1.24	Full residual information protection of resources (FDP_RIP.3)	49
6.1.1.25	Authentication failure handling (FIA_AFL.1)	49
6.1.1.26	User attribute definition (FIA_ATD.1(HU))	49
6.1.1.27	User attribute definition (FIA_ATD.1(TU))	50
6.1.1.28	Verification of secrets (FIA_SOS.1)	50
6.1.1.29	Timing of authentication (FIA_UAU.1)	50
6.1.1.30	Multiple authentication mechanisms (FIA_UAU.5)	50
6.1.1.31	Protected authentication feedback (FIA_UAU.7)	51
6.1.1.32	Timing of identification (FIA_UID.1)	51
6.1.1.33	Enhanced user-subject binding (FIA_USB.2)	51
6.1.1.34	Reliable time stamps (FPT_STM.1)	53
6.1.1.35	Inter-TSF basic TSF data consistency (FPT_TDC.1(BA))	53
6.1.1.36	TSF-initiated session locking (FTA_SSL.1)	53
6.1.1.37	User-initiated locking (FTA_SSL.2)	54

6.1.1.38	Inter-TSF trusted channel (FTP_ITC.1)	54
6.1.2	Virtual machine related functionality	54
6.1.2.1	Complete access control (FDP_ACC.2(VIRT))	54
6.1.2.2	Security attribute based access control (FDP_ACF.1(VIRT))	55
6.1.2.3	Export of user data with security attributes (FDP_ETC.2(VIRT))	55
6.1.2.4	Complete information flow control (FDP_IFC.2(VIRT))	56
6.1.2.5	Simple security attributes (FDP_IFF.1(VIRT))	56
6.1.2.6	Import of user data with security attributes (FDP_ITC.2(VIRT))	57
6.1.2.7	User identification before any action (FIA_UID.2(VIRT))	57
6.1.2.8	Inter-TSF basic TSF data consistency (FPT_TDC.1(VIRT))	57
6.1.3	Management related functionality	58
6.1.3.1	Management of object security attributes (FMT_MSA.1(PSO))	58
6.1.3.2	Management of security attributes (FMT_MSA.1(VIRT-CACP))	58
6.1.3.3	Management of security attributes (FMT_MSA.1(VIRT-CIFCP))	58
6.1.3.4	Static attribute initialisation (FMT_MSA.3(PSO))	58
6.1.3.5	Static attribute initialisation (FMT_MSA.3(NI))	59
6.1.3.6	Static attribute initialisation (FMT_MSA.3(VIRT-CACP))	59
6.1.3.7	Static attribute initialisation (FMT_MSA.3(VIRT-CIFCP))	59
6.1.3.8	Security attribute value inheritance (FMT_MSA.4(PSO))	59
6.1.3.9	Management of TSF data (FMT_MTD.1(AE))	60
6.1.3.10	Management of TSF data (FMT_MTD.1(AS))	60
6.1.3.11	Management of TSF data (FMT_MTD.1(AT))	60
6.1.3.12	Management of TSF data (FMT_MTD.1(AF))	61
6.1.3.13	Management of TSF data (FMT_MTD.1(NI))	61
6.1.3.14	Management of TSF data (FMT_MTD.1(IAT))	61
6.1.3.15	Management of TSF data (FMT_MTD.1(IAF))	61
6.1.3.16	Management of TSF data (FMT_MTD.1(IAU))	61
6.1.3.17	Management of TSF data (FMT_MTD.1(SSH))	62
6.1.3.18	Management of TSF data (FMT_MTD.1(SSL))	62
6.1.3.19	Management of TSF data (FMT_MTD.1(VIRT-COMP))	62
6.1.3.20	Revocation (FMT_REV.1(OBJ))	62
6.1.3.21	Revocation (FMT_REV.1(USR))	63
6.1.3.22	Specification of management functions (FMT_SMF.1)	63
6.1.3.23	Security roles (FMT_SMR.1)	63
6.2	Security Functional Requirements Rationale	64
6.2.1	Coverage	64
6.2.2	Sufficiency	67
6.2.3	Security requirements dependency analysis	69
6.3	Security Assurance Requirements	73
6.4	Security Assurance Requirements Rationale	75
7	TOE Summary Specification	76
7.1	TOE Security Functionality	76
7.1.1	Audit	76
7.1.1.1	Audit functionality	76

7.1.1.2	Audit trail	77
7.1.2	Cryptographic services	78
7.1.2.1	SSHv2 Protocol	78
7.1.3	Packet filter	79
7.1.3.1	Network layer filtering	79
7.1.3.2	Link layer filtering	81
7.1.4	Identification and Authentication	82
7.1.4.1	PAM-based identification and authentication mechanisms	82
7.1.4.2	User Identity Changing	83
7.1.4.3	Authentication Data Management	84
7.1.4.4	SSH key-based authentication	84
7.1.4.5	Session locking	85
7.1.5	Discretionary Access Control	85
7.1.5.1	Permission bits	86
7.1.5.2	Access Control Lists (ACLs)	86
7.1.5.3	File system objects	86
7.1.6	Authoritative Access Control	87
7.1.6.1	Resource access control for virtual machines	87
7.1.7	Virtual machine environments	87
7.1.8	Security Management	90
7.1.8.1	Privileges	90
8	Abbreviations, Terminology and References	91
8.1	Abbreviations	91
8.2	Terminology	91
8.3	References	94

List of Tables

Table 1: Non-evaluated functionalities	15
Table 2: Mapping of security objectives to threats and policies	25
Table 3: Mapping of security objectives for the Operational Environment to assumptions, threats and policies	26
Table 4: Sufficiency of objectives countering threats	27
Table 5: Sufficiency of objectives holding assumptions	29
Table 6: Sufficiency of objectives enforcing Organizational Security Policies	30
Table 7: Security functional requirements for the TOE	35
Table 8: Mapping of security functional requirements to security objectives	64
Table 9: Security objectives for the TOE rationale	67
Table 10: TOE SFR dependency analysis	69
Table 11: Security assurance requirements	74
Table 12: SSH implementation notes	78

List of Figures

Figure 1: KVM virtual machine handling by the Linux kernel	88
--	----

1 Introduction

1.1 Security Target Identification

Title: Red Hat Enterprise Linux 5.6 KVM Security Target
Version: 0.12
Status: Released
Date: 2011-05-18
Sponsor: Red Hat, Inc.
Developer: Red Hat, Inc.
Certification ID: BSI-DSZ-CC-0724
Keywords: Security Target, Common Criteria, Linux Distribution, Embedded Linux

1.2 TOE Identification

The TOE is Red Hat Enterprise Linux Version 5.6.

1.3 TOE Type

The TOE type is Linux-based general-purpose operating system. The evaluation focuses on the virtual machine environment provided with KVM.

1.4 TOE Overview

This security target documents the security characteristics of the Red Hat Enterprise Linux distribution (abbreviated with RHEL throughout this document).

Red Hat Enterprise Linux is a highly-configurable Linux-based operating system which has been developed to provide a good level of security as required in commercial environments. This evaluation focuses on the use of RHEL to provide a KVM virtualization environment. RHEL provides the host system for the virtual machine environment and manages the virtual machines using the KVM technology. In addition, RHEL provides management interfaces to administer the virtual machine environment as well as full auditing of user and administrator operations.

The KVM technology separates the runtime environment of virtual machines from each other. The Linux kernel operates as the hypervisor to the virtual machines but provides a normal computing environment to administrators of the virtual machines. Therefore, the Linux kernel supports the concurrent execution of virtual machines and regular applications. RHEL uses the processor virtualization support to ensure that the virtual machines execute close to the native speed of the hardware.

In addition to the separation of the runtime environment, RHEL also provides system-inherent separation mechanisms to the resources of virtual machines. This separation ensures that the large software component used for virtualizing and simulating devices executing for each virtual machine cannot interfere with each other. Using the SELinux multi-category mechanism, the virtualization and simulation software instances are isolated. The virtual machine management framework configures SELinux multi-category settings transparently to the administrator.

1.4.1 Required Hardware and Software

The hardware / firmware component of the TOE which allows the installation of the operating system on the following hardware systems:

- IBM System x based on x86 64bit Intel Xeon processors: x3400 M2, x3400 M3, x3500 M2, x3500 M3, x3550 M2, x3550 M3, x3620 M3, x3630 M3, x3650 M2, x3650 M3
- IBM BladeCenter: HS22 and HS22V
- IBM iDataPlex: dx360 M2, dx360 M3

1.4.2 Intended Method of Use

The TOE is a Linux-based multi-user multi-tasking operating system. The TOE may provide services to several users at the same time. After successful login, the users have access to a general computing environment, allowing the start-up of user applications, issuing user commands at shell level, creating and accessing files. The TOE provides adequate mechanisms to separate the users and protect their data. Privileged commands are restricted to administrative users.

The TOE is intended to operate in a networked environment with other instantiations of the TOE as well as other well-behaved peer systems operating within the same management domain. All those systems need to be configured in accordance with a defined common security policy.

It is assumed that responsibility for the safeguarding of the data protected by the TOE can be delegated to human users of the TOE if such users are allowed to log on and spawn processes on their behalf. All data is under the control of the TOE. The data is stored in named objects, and the TOE can associate a description of the access rights to that object with each named object.

The TOE enforces controls such that access to data objects can only take place in accordance with the access restrictions placed on that object by its owner, and by administrative users. Ownership of named objects may be transferred under the control of the access control policies implemented by RHEL.

Discretionary access rights (e.g. read, write, execute) can be assigned to data objects with respect to subjects identified with their UID. Once a subject is granted access to an object, the content of that object may be used freely to influence other objects accessible to this subject.

1.4.2.1 KVM virtual machines

The TOE provides virtual machine environments which allow other operating systems to execute concurrently with the RHEL host system. Each virtual machine is represented as a process in the RHEL host system and is subject to the standard Linux constraints for processes. The virtualization and simulation logic that supports the operation of a virtual machine executes within the same process of the respective virtual machine.

Treating each virtual machine as a normal process by the RHEL host system allows the concurrent execution of standard applications at the same time. Administrative tools are implemented as standard Linux applications. In addition, the computing environment provided by the RHEL host system to users logging into that system does not differ from a standard RHEL system even while virtual machines execute. Therefore, standard Linux applications and daemons may operate concurrently with virtual machines. To access the RHEL host system, the TOE provides console access via OpenSSH as well as access via the virtual machine management daemon provided by libvirt via OpenSSH. In addition, the virtual machine consoles can be accessed using VNC using an

already established OpenSSH channel. The use of OpenSSH for accessing the libvirt management daemon as well as VNC is encapsulated in different client applications which are not part of the TOE, like virsh and virt-manager.

The administration interfaces provided by RHEL ensure that the execution environment as well as the resources allocated to a specific virtual machine are separated from other virtual machine instances. Such a transparent setup aids the administrator in keeping a secure configuration for the host system as well as the virtual machines. Unique SELinux categories are assigned to each virtual machine and its resources. The SELinux policy prevents any access request by a virtual machine to a resource if the SELinux category does not match. In addition, each virtual machine executes with a non-root UID ensuring that the virtual machine operation cannot have an effect on the system policy configured on and enforced by the RHEL host system.

When using RHEL as a host system for virtual machines, normal users are typically not allowed to log on to the system but different UIDs are used to separate different services or virtual machines provided by the system. In such a case, it is assumed that these processes are responsible for the safeguarding of their data. Note: the evaluated configuration permits the use of the host system for regular operation by normal users.

1.4.2.2 Operating environment

The TOE permits one or more processors and attached peripheral and storage devices to be used by multiple applications assigned to different UIDs to perform a variety of functions requiring controlled shared access to the data stored on the system. With different UIDs proper access restrictions to resources assigned to processes can be enforced using the access control mechanisms provided by the TOE. Such installations and usage scenarios are typical for systems accessed by processes or users local to, or with otherwise protected access to, the computer system.

Note: The TOE provides the platform for installing and running arbitrary services. These additional services are not part of the TOE. The TOE is solely the operating system which provides the runtime environment for such services.

All human users, if existent, as well as all services offered by RHEL are assigned unique user identifiers within the single host system that forms the TOE. This user identifier is used together with the attributes and roles assigned to the user identifier as the basis for access control decisions. Except for virtual machine accesses, the TOE authenticates the claimed identity of the user before allowing the user to perform any further actions. Services may be spawned by the TOE without the need for user-interaction. The TOE internally maintains a set of identifiers associated with processes, which are derived from the unique user identifier upon login of the user or from the configured user identifier for a TOE-spawned service. Some of those identifiers may change during the execution of the process according to a policy implemented by the TOE.

1.4.3 Major Security Features

The primary security features of the TOE are specified as part of the **logical boundary description**.

These primary security features are supported by domain separation and reference mediation, which ensure that the features are always invoked and cannot be bypassed.

1.5 TOE Description

1.5.1 Introduction

Red Hat Enterprise Linux is a general purpose, multi-user, multi-tasking Linux based operating system. The evaluated functionality targets the virtual machine functionality provided by the KVM technology. It provides a platform for a variety of applications. In addition, the virtual machine environment provides an execution environment for a large number of different operating systems.

The SELinux security module is configured to enforce the authoritative access control policy. The following access control rules are enforced by SELinux:

- Isolation of virtual machines from each other by assigning each process implementing a virtual machine and its resources a unique category. Access between virtual machines and resources is only permitted if the categories of the virtual machine and the accessed resource are identical.

The RHEL evaluation covers a potentially distributed network of systems running the evaluated versions and configurations of RHEL as well as other peer systems operating within the same management domain. The hardware platforms selected for the evaluation consist of machines which are available when the evaluation has completed and to remain available for a substantial period of time afterwards.

The TOE Security Functions (TSF) consist of functions of RHEL that run in kernel mode plus some trusted processes. These are the functions that enforce the security policy as defined in this Security Target. Tools and commands executed in user mode that are used by an administrative user need also to be trusted to manage the system in a secure way. But as with other operating system evaluations they are not considered to be part of this TSF.

The hardware, the BootProm firmware and potentially other firmware layers between the hardware and the TOE are considered to be part of the TOE environment.

The TOE includes standard networking applications, such as SSH.

System administration tools include the standard command line tools. A graphical user interface for system administration or any other operation is not included in the evaluated configuration.

In addition to providing a general-purpose computing environment, the TOE provides the following mechanisms allowing a wider deployment:

- Virtualization: The KVM technology allows the Linux kernel to act as a hypervisor and execute virtual machines concurrently with standard applications.

The TOE environment also includes applications that are not evaluated, but are used as unprivileged tools to access public system services. For example a network server using a port above 1024 may be used as a normal application running without root privileges on top of the TOE. The additional documentation specific for the evaluated configuration provides guidance how to set up such applications on the TOE in a secure way.

1.5.2 TOE boundaries

1.5.2.1 Physical

The Target of Evaluation is based on the following system software:

- Red Hat Enterprise Linux in the above mentioned version

The TOE and its documentation are supplied on ISO images distributed via the Red Hat Network. The TOE includes a package holding the additional user and administrator documentation.

In addition to the installation media, the following documentation is provided:

- Evaluated Configuration Guide
- Manual pages for all applications, configuration files and system calls

The hardware applicable to the evaluated configuration is listed above. The analysis of the hardware capabilities as well as the firmware functionality is covered by this evaluation to the extent that the following capabilities supporting the security functionality are analyzed and tested:

- Memory separation capability
- Unavailability of privileged processor states to untrusted user code (like the hypervisor state or the SMM)
- Full testing of the security functionality on all listed hardware systems

1.5.2.2 Logical

The primary security features of the TOE are:

- **Auditing:** The Lightweight Audit Framework (LAF) is designed to be an audit system for Linux compliant with the requirements from Common Criteria. LAF is able to intercept all system calls as well as retrieving audit log entries from privileged user space applications. The subsystem allows configuring the events to be actually audited from the set of all events that are possible to be audited.
- **Cryptographically secured communication channels:** The TOE provides cryptographically secured communication to allow remote entities to log into the TOE. For interactive usage, the SSHv2 protocol is provided. In addition, the access to the virtual machine administration tool as well as the virtual machine consoles can be protected using a SSHv2-based tunnel.
- **Packet filter:** The TOE provides a stateless and stateful packet filter for regular IP-based communication. Layer 2 (IP) and layer 3 (TCP, UDP, ICMP) network protocols can be controlled using this packet filter. To allow virtual machines to communicate with the environment, the TOE provides a bridging functionality. Ethernet frames routed through bridges are controlled by a separate packet filter which implements a stateless packet filter for the TCP/IP protocol family as well as VLAN filtering.
- **Identification and Authentication:** User identification and authentication in the TOE includes all forms of interactive login (e.g. using the SSH protocol or log in at the local console) as well as identity changes through the su command. These all rely on explicit authentication information provided interactively by a user.
- **Discretionary Access Control:** DAC allows owners of named objects to control the access permissions to these objects. These owners can permit or deny access for other users based on the configured permission settings. The DAC mechanism is also used to ensure that untrusted users cannot tamper with the TOE mechanisms.
- **Authoritative Access Control:** The TOE supports authoritative access control based on the following concepts:
 - SELinux categories are attached to virtual machines and its resources. The access control policy enforced using these categories grant virtual machines access to resources if the category of the virtual machine is identical to the category of the accessed resource.

Users cannot interfere with these labels. The TOE uses SELinux with an appropriate SELinux policy to enforce the authoritative access control.

- Virtual machine environments: The TOE implements the host system for virtual machines. It acts as a hypervisor which provides an environment to allow other operating systems execute concurrently.
- Security Management: The security management facilities provided by the TOE are usable by authorized users and/or authorized administrators to modify the configuration of TSF.

The TOE provides many more functions and mechanisms. The evaluation ensures that all these additional functions do not interfere with the above mentioned security mechanisms in the evaluated configuration. Mechanisms and functions that would interfere with the operation of the security functions are disallowed in the evaluated configuration and the Evaluation Configuration Guide provides instructions to the administrator on how to disable them. Note: TOE mechanism which provide additional restrictions to the above claimed security functions are allowed in the evaluated configuration. For example, the SELinux-based MLS and Type Enforcement policy is provided and enforced with the TOE and permitted in the evaluated configuration even though they have not been subject to this evaluation. The MLS and Type Enforcement policy provides further restrictions on, for example, the security function of discretionary access control mechanism for IPC objects and therefore cannot breach the security functionality. The following table enumerates mechanisms that are provided with the TOE but which are excluded from the evaluation:

Functions	Exclusion discussion
eCryptFS, dm-crypt	eCryptFS and dm-crypt are not allowed to be used in the evaluated configuration. The encryption capability provided with this file system and the device mapper are therefore unavailable to any user.
SSL / TLS tunnels	The TOE provides the stunnel application which can be used to establish SSL and TLS tunnels with remote peers. This application however was excluded from evaluation assessment.
SELinux MLS and type enforcement	The MLS and type enforcement policy provided with SELinux is not subject to evaluation.

Table 1: Non-evaluated functionalities

Note: Packages and mechanisms not covered with security claims and subsequent assessments during the evaluation or disabling the respective functionality in the evaluated configuration result from resource constraints during the evaluation but does not imply that the respective package or functionality is implemented insecurely.

1.5.2.3 Configurations

The evaluated configurations are defined as follows:

- The CC evaluated package set must be selected at install time in accordance with the description provided in the Evaluated Configuration Guide and installed accordingly.
- The TOE supports the use of IPv4 and IPv6, both are also supported in the evaluated configuration.

- The default configuration for identification and authentication are the defined password-based PAM modules as well as by the certificate based authentication for OpenSSH. Support for other authentication options, e.g. smart card authentication, is not included in the evaluation configuration.
- If the system console is used, it must be connected directly to the TOE and afforded the same physical protection as the TOE.

Deviations from the configurations and settings specified with the Evaluated Configuration Guide are not permitted.

The TOE comprises a single system (and optional peripherals) running the TOE software listed. Cluster configurations are not permitted in the evaluated configuration.

1.5.2.4 TOE Environment

Several TOE systems may be interlinked in a network, and individual networks may be joined by bridges and/or routers, or by TOE systems which act as routers and/or gateways. Each of the TOE systems implements its own security policy. The TOE does not include any synchronization function for those policies. As a result a single user may have user accounts on each of those systems with different UIDs, different roles, and other different attributes. (A synchronization method may optionally be used, but it not part of the TOE and must not use methods that conflict with the TOE requirements.)

If other systems are connected to a network they need to be configured and managed by the same authority using an appropriate security policy that does not conflict with the security policy of the TOE. All links between this network and untrusted networks (e. g. the Internet) need to be protected by appropriate measures such as carefully configured firewall systems that prohibit attacks from the untrusted networks. Those protections are part of the TOE environment.

1.5.2.5 Security Policy Model

The security policy for the TOE is defined by the security functional requirements in chapter 6. The following is a list of the subjects and objects participating in the policy.

Subjects:

- Processes acting on behalf of a human user or technical entity.
- Processes acting on behalf of a human user or technical entity providing a virtual machine environment.

Named objects:

- File system objects in the following allowed file systems:
 - Ext3 - standard file system for general data
 - iso9660 - ISO9660 file system for CD-ROM and DVD
 - tmpfs - the temporary file system backed by RAM
 - rootfs - the virtual root file system used temporarily during system boot
 - procfs - process file system holding information about processes, general statistical data and tunable kernel parameters
 - sysfs - system-related file system covering general information about resources maintained by the kernel including several tunable parameters for these resources
 - devpts - pseudoterminal file system for allocating virtual TTYs on demand

- binfmt_misc - configuration interface allowing the assignment of executable file formats with user space applications
- securityfs - interface for loadable security modules (LSM) to provide tunables and configuration interfaces to user space
- selinuxfs - interface for allowing user space components to interact with the SELinux module inside the kernel, including managing the SELinux policy.

Please note that the TOE supports a number of additional virtual (i.e. without backing of persistent storage) file systems which are only accessible to the TSF - they are not or cannot be mounted. All above mentioned virtual file systems implement access decisions based DAC attributes inferred from the underlying process' DAC attributes. Additional restrictions may apply for specific objects in this file system.

- Virtual machine resources:
 - Disks: files as backing store
 - Disks: device files as backing store
 - Network device: any physical or logical network device present on the host system
 - Non-network devices: any device present on the host system, including USB and bluetooth

Please note that CPU restrictions and memory range assignments are not listed as part of virtual machine resources, because virtual machines cannot instruct the host system to access these resources. The host system scheduler automatically enforces CPU restrictions. Also, the host system virtual memory management functionality automatically enforces the memory range restriction.

TSF data:

- Subject meta data - all data used for subjects except data which is not interpreted by the TSF and does not implement parts of the TSF (this data is called user data)
- Named object meta data - all data used for the respective objects except data which is not interpreted by the TSF and does not implement parts of the TSF (this data is called user data)
- User accounts, including the security attributes defined by FIA_ATD.1
- Audit records

User data:

- Non-TSF executable code used to drive the behavior of subjects
- Data not interpreted by TSF and stored or transmitted using named objects
- Any code executed within the virtual machine environment as well as any data stored in resources assigned to virtual machines

2 CC Conformance Claim

This ST is CC Part 2 extended and CC Part 3 conformant, with a claimed Evaluation Assurance Level of EAL4, augmented by ALC_FLR.3.

This ST does not claim conformance to any Protection Profile.

Common Criteria [CC] version 3.1 revision 3 is the basis for this conformance claim.

3 Security Problem Definition

3.1 Threat Environment

Threats to be countered by the TOE are characterized by the combination of an asset being subject to a threat, a threat agent and an adverse action.

The definition of threat agents and protected assets that follows is applicable to the OSPP base, as well as to the OSPP extended packages, unless noted otherwise.

3.1.1 Assets

Assets to be protected are:

- Persistent storage objects used to store user data and/or TSF data, where this data needs to be protected from any of the following operations:
 - Unauthorized read access
 - Unauthorized modification
 - Unauthorized deletion of the object
 - Unauthorized creation of new objects
 - Unauthorized management of object attributes
- TSF functions and associated TSF data
- The resources managed by the TSF that are used to store the above-mentioned objects, including the metadata needed to manage these objects.

3.1.2 Threat Agents

Threat agents are external entities that potentially may attack the TOE. They satisfy one or more of the following criteria:

- External entities not authorized to access assets may attempt to access them either by masquerading as an authorized entity or by attempting to use TSF services without proper authorization.
- External entities authorized to access certain assets may attempt to access other assets they are not authorized to either by misusing services they are allowed to use or by masquerading as a different external entity.
- Untrusted subjects may attempt to access assets they are not authorized to either by misusing services they are allowed to use or by masquerading as a different subject.

Threat agents are typically characterized by a number of factors, such as expertise, available resources, and motivation, with motivation being linked directly to the value of the assets at stake. The TOE protects against intentional and unintentional breach of TOE security by attackers possessing an enhanced-basic attack potential.

3.1.3 Threats countered by the TOE

T.ACCESS.TSFDATA

A threat agent might read or modify TSF data without the necessary authorization when the data is stored or transmitted.

T.ACCESS.USERDATA

A threat agent might gain access to user data stored, processed or transmitted by the TOE without being appropriately authorized according to the TOE security policy.

T.ACCESS.TSFFUNC

A threat agent might use or modify functionality of the TSF without the necessary privilege to grant itself or others unauthorized access to TSF data or user data.

T.ACCESS.COMM

A threat agent might access a communication channel that establishes a trust relationship between the TOE and another remote trusted IT system or masquerade as another remote trusted IT system.

T.RESTRICT.NETTRAFFIC

A threat agent might get access to information or transmit information to other recipients via network communication channels without authorization for this communication attempt by the information flow control policy.

T.IA.MASQUERADE

A threat agent might masquerade as an authorized entity including the TOE itself or a part of the TOE in order to gain unauthorized access to user data, TSF data, or TOE resources.

T.IA.USER

A threat agent might gain access to user data, TSF data or TOE resources with the exception of public objects without being identified and authenticated.

T.ACCESS.COMPENV

A threat agent might utilize or modify the runtime environment of other compartments in an unauthorized manner.

T.INFOFLOW.COMP

A threat agent might get access to information without authorization by the information flow control policy.

T.COMM.COMP

A threat agent might access the data communicated between compartments or between a compartment and an external entity to read or modify the transferred data.

3.2 Assumptions

3.2.1 Environment of use of the TOE

3.2.1.1 Physical

A.PHYSICAL

It is assumed that the IT environment provides the TOE with appropriate physical security, commensurate with the value of the IT assets protected by the TOE.

3.2.1.2 Personnel

A.MANAGE

The TOE security functionality is managed by one or more competent individuals. The system administrative personnel are not careless, willfully negligent, or hostile, and will follow and abide by the instructions provided by the guidance documentation.

A.AUTHUSER

Authorized users possess the necessary authorization to access at least some of the information managed by the TOE and are expected to act in a cooperating manner in a benign environment.

A.TRAINEDUSER

Users are sufficiently trained and trusted to accomplish some task or group of tasks within a secure IT environment by exercising complete control over their user data.

3.2.1.3 Procedural

A.DETECT

Any modification or corruption of security-enforcing or security-relevant files of the TOE, user or the underlying platform caused either intentionally or accidentally will be detected by an administrative user.

A.PEER.MGT

All remote trusted IT systems trusted by the TSF to provide TSF data or services to the TOE, or to support the TSF in the enforcement of security policy decisions are assumed to be under the same management control and operate under security policy constraints compatible with those of the TOE.

A.PEER.FUNC

All remote trusted IT systems trusted by the TSF to provide TSF data or services to the TOE, or to support the TSF in the enforcement of security policy decisions are assumed to correctly implement the functionality used by the TSF consistent with the assumptions defined for this functionality.

3.2.1.4 Connectivity

A.CONNECT

All connections to and from remote trusted IT systems and between physically-separate parts of the TSF not protected by the TSF itself are physically or logically protected within the TOE environment to ensure the integrity and confidentiality of the data transmitted and to ensure the authenticity of the communication end points.

3.3 Organizational Security Policies

P.ACCOUNTABILITY

The users of the TOE shall be held accountable for their security-relevant actions within the TOE.

P.USER

Authority shall only be given to users who are trusted to perform the actions correctly.

P.PROTECT_SSH_KEY

When using SSH with key-based authentication, organizational procedures must exist that ensure users protect their private SSH key component against its use by any other user.

Note: The protection of the key can be established by access permissions to the file holding the key (when using the OpenSSH client, the key file permissions are automatically verified and the key is rejected if the permissions are not restrictive), or by encrypting the key with a passphrase. Making the SSH private key available to any other user is akin to telling that user the password for password-based authentication.

4 Security Objectives

4.1 Objectives for the TOE

O.AUDITING

The TSF must be able to record defined security-relevant events (which usually include security-critical actions of users of the TOE). The TSF must protect this information and present it to authorized users if the audit trail is stored on the local system. The information recorded for security-relevant events must contain the time and date the event happened and, if possible, the identification of the user that caused the event, and must be in sufficient detail to help the authorized user detect attempted security violations or potential misconfiguration of the TOE security features that would leave the IT assets open to compromise.

O.CRYPTO.NET

The TSF must allow authorized users to remotely access the TOE using a cryptographically-protected network protocol that ensures integrity and confidentiality of the transported data and is able to authenticate the end points of the communication. Note that the same protocols may also be used in the case where the TSF is physically separated into multiple parts that must communicate securely with each other over untrusted network connections.

O.DISCRETIONARY.ACCESS

The TSF must control access of subjects and/or users to named resources based on identity of the object. The TSF must allow authorized users to specify for each access mode which users/subjects are allowed to access a specific named object in that access mode.

O.NETWORK.FLOW

The TOE shall mediate communication between sets of TOE network interfaces, between a network interface and the TOE itself, and between subjects in the TOE and the TOE itself in accordance with its security policy.

O.I&A

The TOE must ensure that users have been successfully authenticated before allowing any action the TOE has defined to provide to authenticated users only.

O.MANAGE

The TSF must provide all the functions and facilities necessary to support the authorized users that are responsible for the management of TOE security mechanisms and must ensure that only authorized users are able to access such functionality.

O.TRUSTED_CHANNEL

The TSF must be designed and implemented in a manner that allows for establishing a trusted channel between the TOE and a remote trusted IT system that protects the user data and TSF data transferred over this channel from disclosure and undetected modification and prevents masquerading of the remote trusted IT system.

O.COMP.INFO_FLOW_CTRL

The TOE will control information flow between compartments under the control of the TOE, based on security attributes of these compartments and potentially other TSF data (e.g., security attributes of objects). This information flow control policy must be able to allow the isolation of individual compartments from other compartments controlled by the TOE.

O.COMP.RESOURCE_ACCESS

The TOE will control access of compartments to objects and resources under its control based on:

- security attributes of the objects,
- security attributes of the compartment that attempts to access the object, and
- the type of access attempted.

The rules that determine access may be based on the value of other TSF data. Access must be controlled down to individual compartments and objects.

O.COMP.IDENT

For each access request, the TOE is able to identify the compartment requesting to access resources, objects or information.

4.2 Objectives for the Operational Environment

OE.ADMIN

Those responsible for the TOE are competent and trustworthy individuals, capable of managing the TOE and the security of the information it contains.

OE.REMOTE

If the TOE relies on remote trusted IT systems to support the enforcement of its policy, those systems provide the functions required by the TOE and are sufficiently protected from any attack that may cause those functions to provide false results.

OE.INFO_PROTECT

Those responsible for the TOE must establish and implement procedures to ensure that information is protected in an appropriate manner. In particular:

- All network and peripheral cabling must be approved for the transmittal of the most sensitive data held by the system. Such physical links are assumed to be adequately protected against threats to the confidentiality and integrity of the data transmitted.
- DAC protections on security-relevant files (such as audit trails and authentication databases) shall always be set up correctly.
- Users are authorized to access parts of the data managed by the TOE and are trained to exercise control over their own data.

OE.INSTALL

Those responsible for the TOE must establish and implement procedures to ensure that the hardware, software and firmware components that comprise the system are distributed, installed and configured in a secure manner supporting the security mechanisms provided by the TOE.

OE.MAINTENANCE

Authorized users of the TOE must ensure that the comprehensive diagnostics facilities provided by the product are invoked at every scheduled preventative maintenance period.

OE.PHYSICAL

Those responsible for the TOE must ensure that those parts of the TOE critical to enforcement of the security policy are protected from physical attack that might compromise IT security objectives. The protection must be commensurate with the value of the IT assets protected by the TOE.

OE.RECOVER

Those responsible for the TOE must ensure that procedures and/or mechanisms are provided to assure that after system failure or other discontinuity, recovery without a protection (security) compromise is achieved.

OE.TRUSTED.IT.SYSTEM

The remote trusted IT systems implement the protocols and mechanisms required by the TSF to support the enforcement of the security policy.

These remote trusted IT systems are under the same management domain as the TOE, are managed based on the same rules and policies applicable to the TOE, and are physically and logically protected equivalent to the TOE.

4.3 Security Objectives Rationale

4.3.1 Coverage

The following table provides a mapping of TOE objectives to threats and policies, showing that each objective counters or enforces at least one threat or policy, respectively.

Objective	Threats / OSPs
O.AUDITING	P.ACCOUNTABILITY
O.CRYPTO.NET	T.ACCESS.TSFDATA T.ACCESS.USERDATA T.ACCESS.TSFFUNC
O.DISCRETIONARY.ACCESS	T.ACCESS.TSFDATA T.ACCESS.USERDATA
O.NETWORK.FLOW	T.RESTRICT.NETTRAFFIC
O.I&A	T.IA.MASQUERADE T.IA.USER
O.MANAGE	T.ACCESS.TSFFUNC P.ACCOUNTABILITY P.USER
O.TRUSTED_CHANNEL	T.ACCESS.COMM
O.COMP.INFO_FLOW_CTRL	T.INFOFLOW.COMP

Objective	Threats / OSPs
O.COMP.RESOURCE_ACCESS	T.ACCESS.COMPENV T.COMM.COMP
O.COMP.IDENT	T.ACCESS.COMPENV T.INFOFLOW.COMP T.COMM.COMP

Table 2: Mapping of security objectives to threats and policies

The following table provides a mapping of the objectives for the Operational Environment to assumptions, threats and policies, showing that each objective holds, counters or enforces at least one assumption, threat or policy, respectively.

Objective	Assumptions / Threats / OSPs
OE.ADMIN	A.MANAGE A.AUTHUSER A.TRAINEDUSER
OE.REMOTE	A.CONNECT T.ACCESS.COMM
OE.INFO_PROTECT	A.PHYSICAL A.MANAGE A.AUTHUSER A.TRAINEDUSER P.USER P.PROTECT_SSH_KEY
OE.INSTALL	A.MANAGE A.DETECT
OE.MAINTENANCE	A.DETECT
OE.PHYSICAL	A.PHYSICAL
OE.RECOVER	A.MANAGE A.DETECT
OE.TRUSTED.IT.SYSTEM	A.PEER.MGT A.PEER.FUNC A.CONNECT

Table 3: Mapping of security objectives for the Operational Environment to assumptions, threats and policies

4.3.2 Sufficiency

The following rationale provides justification that the security objectives are suitable to counter each individual threat and that each security objective tracing back to a threat, when achieved, actually contributes to the removal, diminishing or mitigation of that threat:

Threat	Rationale for security objectives
T.ACCESS.TSFDATA	<p>The threat of accessing TSF data without proper authorization is removed by:</p> <ul style="list-style-type: none"> • O.CRYPTO.NET requiring cryptographically-protected communication channels for data including TSF data controlled by the TOE in transit between trusted IT systems. • O.DISCRETIONARY.ACCESS requiring that data, including TSF data stored with the TOE, have discretionary access control protection.
T.ACCESS.USERDATA	<p>The threat of accessing user data without proper authorization is removed by:</p> <ul style="list-style-type: none"> • O.CRYPTO.NET requiring cryptographically-protected communication channels for data including user data controlled by the TOE in transit between trusted IT systems. • O.DISCRETIONARY.ACCESS requiring that data including user data stored with the TOE, have discretionary access control protection.
T.ACCESS.TSFFUNC	<p>The threat of accessing TSF functions without proper authorization is removed by:</p> <ul style="list-style-type: none"> • O.CRYPTO.NET requiring cryptographically-protected communication channels to limit which TSF functions are accessible to external entities. • O.MANAGE requiring that only authorized users utilize management TSF functions.
T.ACCESS.COMM	<p>The threat of accessing a communication channel that establishes a trust relationship between the TOE and another remote trusted IT system is removed by:</p> <ul style="list-style-type: none"> • O.TRUSTED_CHANNEL requiring that the TOE implements a trusted channel between itself and a remote trusted IT system protecting the user data and TSF data transferred over this channel from disclosure and undetected modification and prevents masquerading of the remote trusted IT system. • OE.REMOTE requiring that those systems providing the functions required by the TOE are sufficiently protected from any attack that may cause those functions to provide false results.
T.RESTRICT.NETTRAFFIC	<p>The threat of accessing information or transmitting information to other recipients via network communication channels without authorization for this communication attempt is removed by:</p> <ul style="list-style-type: none"> • O.NETWORK.FLOW requiring the TOE to mediate the communication between itself and remote entities in accordance with its security policy.
T.IA.MASQUERADE	<p>The threat of masquerading as an authorized entity in order to gain unauthorized access to user data, TSF data or TOE resources is removed by:</p>

Threat	Rationale for security objectives
	<ul style="list-style-type: none"> • O.I&A requiring that each entity interacting with the TOE is properly identified and authenticated before allowing any action the TOE is defined to provide to authenticated users only.
T.IA.USER	<p>The threat of accessing user data, TSF data or TOE resources without being identified and authenticated is removed by:</p> <ul style="list-style-type: none"> • O.I&A requiring that each entity interacting with the TOE is properly identified and authenticated before allowing any action the TOE has defined to provide to authenticated users only.
T.ACCESS.COMPENV	<p>The threat of utilizing or modifying the runtime environment of compartments executing on behalf of other users is removed by:</p> <ul style="list-style-type: none"> • O.COMP.RESOURCE_ACCESS requiring the TOE to control access of compartments to objects and resources under its control. • O.COMP.IDENT requiring the TOE to identify the compartment requesting to access resources, objects or information for each access request.
T.INFOFLOW.COMP	<p>The threat of accessing information without authorization by the information flow control policy is removed by:</p> <ul style="list-style-type: none"> • O.COMP.INFO_FLOW_CTRL requiring the TOE to control information flow between compartments under the control of the TOE based on security attributes of these compartments and potentially other TSF data. • O.COMP.IDENT requiring the TOE to identify the compartment requesting to access resources, objects or information for each access request.
T.COMM.COMP	<p>The threat of accessing the data communicated between compartments or between a compartment and an external entity is removed by:</p> <ul style="list-style-type: none"> • O.COMP.RESOURCE_ACCESS requiring the TOE to control access of compartments to objects and resources under its control. • O.COMP.IDENT requiring the TOE to identify the compartment requesting to access resources, objects or information for each access request.

Table 4: Sufficiency of objectives countering threats

The following rationale provides justification that the security objectives for the environment are suitable to cover each individual assumption, that each security objective for the environment that traces back to an assumption about the environment of use of the TOE, when achieved, actually contributes to the environment achieving consistency with the assumption, and that if all security objectives for the environment that trace back to an assumption are achieved, the intended usage is supported:

Assumption	Rationale for security objectives
A.PHYSICAL	<p>The assumption on the IT environment to provide the TOE with appropriate physical security, commensurate with the value of the IT assets protected by the TOE is covered by:</p> <ul style="list-style-type: none"> • OE.INFO_PROTECT requiring the approval of network and peripheral cabling. • OE.PHYSICAL requiring physical protection.
A.MANAGE	<p>The assumptions on the TOE security functionality being managed by one or more trustworthy individuals is covered by:</p> <ul style="list-style-type: none"> • OE.ADMIN requiring trustworthy personnel managing the TOE. • OE.INFO_PROTECT requiring personnel to ensure that information is protected in an appropriate manner. • OE.INSTALL requiring personnel to ensure that components that comprise the system are distributed, installed and configured in a secure manner supporting the security mechanisms provided by the TOE. • OE.RECOVER requiring personnel to assure that after system failure or other discontinuity, recovery without a protection (security) compromise is achieved.
A.AUTHUSER	<p>The assumption on authorized users to possess the necessary authorization to access at least some of the information managed by the TOE and to act in a cooperating manner in a benign environment is covered by:</p> <ul style="list-style-type: none"> • OE.ADMIN ensuring that those responsible for the TOE are competent and trustworthy individuals, capable of managing the TOE and the security of the information it contains. • OE.INFO_PROTECT requiring that DAC protections on security-relevant files (such as audit trails and authentication databases) shall always be set up correctly and that users are authorized to access parts of the data maintained by the TOE.
A.TRAINEDUSER	<p>The assumptions on users to be sufficiently trained and trusted to accomplish some task or group of tasks within a secure IT environment by exercising complete control over their user data is covered by:</p> <ul style="list-style-type: none"> • OE.ADMIN requiring competent personnel managing the TOE. • OE.INFO_PROTECT requiring that those responsible for the TOE must establish and implement procedures to ensure that information is protected in an appropriate manner and that users are trained to exercise control over their own data.
A.DETECT	<p>The assumption that modification or corruption of security-enforcing or security-relevant files will be detected by an administrative user is covered by:</p> <ul style="list-style-type: none"> • OE.INSTALL requiring an administrative user to ensure that the TOE is distributed, installed and configured in a secure manner supporting the security mechanisms provided by the TOE.

Assumption	Rationale for security objectives
	<ul style="list-style-type: none"> • OE.MAINTENANCE requiring an administrative user to ensure that the diagnostics facilities are invoked at every scheduled preventative maintenance period, verifying the correct operation of the TOE. • OE.RECOVER requiring an administrative user to ensure that procedures and/or mechanisms are provided to assure that after system failure or other discontinuity, recovery without a protection (security) compromise is achieved.
A.PEER.MGT	<p>The assumption on all remote trusted IT systems to be under the same management control and operate under security policy constraints compatible with those of the TOE is covered by:</p> <ul style="list-style-type: none"> • OE.TRUSTED.IT.SYSTEM requiring that these remote trusted IT systems are under the same management domain as the TOE, and are managed based on the same rules and policies applicable to the TOE.
A.PEER.FUNC	<p>The assumption on all remote trusted IT systems to correctly implement the functionality used by the TSF consistent with the assumptions defined for this functionality is covered by:</p> <ul style="list-style-type: none"> • OE.TRUSTED.IT.SYSTEM requiring that the remote trusted IT systems implement the protocols and mechanisms required by the TSF to support the enforcement of the security policy.
A.CONNECT	<p>The assumption on all connections to and from remote trusted IT systems and between physically separate parts of the TSF not protected by the TSF itself are physically or logically protected is covered by:</p> <ul style="list-style-type: none"> • OE.REMOTE requiring that remote trusted IT systems provide the functions required by the TOE and are sufficiently protected from any attack that may cause those functions to provide false results. • OE.TRUSTED.IT.SYSTEM demanding the physical and logical protection equivalent to the TOE.

Table 5: Sufficiency of objectives holding assumptions

The following rationale provides justification that the security objectives are suitable to cover each individual organizational security policy, that each security objective that traces back to an OSP, when achieved, actually contributes to the implementation of the OSP, and that if all security objectives that trace back to an OSP are achieved, the OSP is implemented:

OSP	Rationale for security objectives
P.ACCOUNTABILITY	<p>The policy to hold users accountable for their security-relevant actions within the TOE is implemented by:</p> <ul style="list-style-type: none"> • O.AUDITING providing the TOE with audit functionality. • O.MANAGE allowing the management of this function.

OSP	Rationale for security objectives
P.USER	<p>The policy to match the trust given to a user and the actions the user is given authority to perform is implemented by:</p> <ul style="list-style-type: none">• O.MANAGE allowing appropriately-authorized users to manage the TSF.• OE.INFO_PROTECT, which requires that users are trusted to use the protection mechanisms of the TOE to protect their data.
P.PROTECT_SSH_KEY	<p>The policy to match the trust given to a user to protect his SSH private key is implemented by:</p> <ul style="list-style-type: none">• OE.INFO_PROTECT, which requires that users are trusted to exercise the control over their own data.

Table 6: Sufficiency of objectives enforcing Organizational Security Policies

5 Extended Components Definition

The Security Target defines the extended component FCS_RNG as part of the FCS class in CC Part 2 for usage within this ST. In addition, the Security Target defines the extended component FDP_RIP.3 as part of the FDP_RIP family in CC Part 2 for usage within this ST and FIA_USB.2 as part of the FIA_USB family in CC Part 2.

5.1 Class FCS: Cryptographic support

5.1.1 Random number generator (RNG)

Family behaviour

This family defines quality requirements for the generation of random numbers that are intended to be used for cryptographic purposes.

Component levelling

FCS_RNG.1 is not hierarchical to any other component within the FCS_RNG family.

Management: FCS_RNG.1

The following actions could be considered for the management functions in FMT:

- a) There are no management activities foreseen.

Audit: FCS_RNG.1

The following actions should be auditable if FAU_GEN Security audit data generation is included in the PP/ST:

- a) Minimal: There are no actions defined to be auditable.
- b) Basic: There are no actions defined to be auditable.
- c) Detailed: There are no actions defined to be auditable.

5.1.1.1 FCS_RNG.1 - Random number generation

Hierarchical to: No other components.

Dependencies: No dependencies.

FCS_RNG.1.1 The TSF shall provide a [selection: **physical, non-physical true, deterministic, physical hybrid, deterministic hybrid**] random number generator that implements: [assignment: **list of security capabilities**].

FCS_RNG.1.2 The TSF shall provide random numbers that meet [assignment: **a defined quality metric**].

Rationale

The quality of the random number generator is defined using this SFR. The quality metric required in FCS_RNG.1.2 is detailed in the German Scheme AIS 20 and AIS 31.

5.2 Class FDP: User data protection

FDP_RIP.3 is analog to FDP_RIP.2 except that it applies to the content of resources that are allocated to subjects or users.

5.2.1 Residual information protection (RIP)

Component levelling

The SFR is not hierarchical to any SFR out of the family of FDP_RIP.

Management: FDP_RIP.3

The following actions could be considered for the management functions in FMT:

- a) See management description specified for FDP_RIP.2 in [CC] part 2.

Audit: FDP_RIP.3

The following actions should be auditable if FAU_GEN Security audit data generation is included in the PP/ST:

- a) Minimal: See audit requirement specified for FDP_RIP.2 in [CC] part 2.
- b) Basic: See audit requirement specified for FDP_RIP.2 in [CC] part 2.
- c) Detailed: See audit requirement specified for FDP_RIP.2 in [CC] part 2.

5.2.1.1 FDP_RIP.3 - Full residual information protection of resources

Hierarchical to: No other components.

Dependencies: No dependencies.

FDP_RIP.3.1 The TSF shall ensure that any previous information content of a resource is made unavailable upon the [selection: **allocation of the resource to, deallocation of the resource from**] all subjects or users.

Rationale

FDP_RIP.3 addresses the problem of resources implemented in main memory that may be allocated to and de-allocated from subjects or users. Unless those resources lose their content automatically as part of the de-allocation and re-allocation process, they must be subject to a process that prepares them for re-use by rendering the previous content unavailable to the subject or user to which it is next allocated. An example is main memory that has been allocated to a subject; this memory must be cleared before it can be re-allocated to a subject with different security attributes (for example a subject operating on behalf of a different user). This preparation prevents the passing of security-critical information via this resource, since such unregulated passing would potentially allow the subject or user to which the memory is next allocated to use this information to violate the security policy. Typical examples of such critical information that may be passed via resources not prepared for re-use are passwords or cryptographic keys.

5.3 Class FIA: Identification and authentication

FIA_USB.2 is analog to FIA_USB.1 except that it adds the possibility to specify rules whereby subject security attributes are also derived from TSF data other than user security attributes.

5.3.1 User-subject binding (USB)

Component levelling

FIA_USB.2 is hierarchical to FIA_USB.1.

Management: FIA_USB.2

The following actions could be considered for the management functions in FMT:

- a) See management description specified for FIA_USB.1 in [CC] part 2.

Audit: FIA_USB.2

The following actions should be auditable if FAU_GEN Security audit data generation is included in the PP/ST:

- a) Minimal: See audit requirement specified for FIA_USB.1 in [CC] part 2.
- b) Basic: See audit requirement specified for FIA_USB.1 in [CC] part 2.
- c) Detailed: See audit requirement specified for FIA_USB.1 in [CC] part 2.

5.3.1.1 FIA_USB.2 - Enhanced user-subject binding

Hierarchical to: FIA_USB.1 User-subject binding

Dependencies: FIA_ATD.1 User attribute definition

- FIA_USB.2.1** The TSF shall associate the following user security attributes with subjects acting on the behalf of that user: [assignment: **list of user security attributes**].
- FIA_USB.2.2** The TSF shall enforce the following rules on the initial association of user security attributes with subjects acting on the behalf of users: [assignment: **rules for the initial association of attributes**].
- FIA_USB.2.3** The TSF shall enforce the following rules governing changes to the user security attributes associated with subjects acting on the behalf of users: [assignment: **rules for the changing of attributes**].
- FIA_USB.2.4** The TSF shall enforce the following rules for the assignment of subject security attributes not derived from user security attributes when a subject is created: [assignment: **rules for the initial association of the subject security attributes not derived from user security attributes**].

Rationale

An operating system may derive subject security attributes from other TSF data that are not directly user security attributes. An example is the point-of-entry the user has used to establish the connection. An access control policy may also use this subject security attribute within its access control policy, allowing access to critical objects only when the user has connected through specific ports-of-entry.

6 Security Requirements

6.1 TOE Security Functional Requirements

All of the following SFRs are derived from the OSPP and the OSPP extended package on virtualization. Although this ST does not claim any compliance to the OSPP (as not all SFRs of the OSPP are included), the SFRs taken from OSPP are instantiated exactly as the OSPP would require it. Therefore, some SFRs may contain spurious "no other rule" or similar assignments which would make no sense if the SFRs are derived from CC part 2. As this ST will be reused in a later evaluation to be fully OSPP compliant, the ST author decided to already apply all OSPP requirements to the instantiated SFRs.

The following table shows the Security functional requirements for the TOE, and the operations performed on the components according to CC part 2: iteration (Iter.), refinement (Ref.), assignment (Ass.) and selection (Sel.).

Security functional group	Security functional requirement	Base security functional component	Source	Operations			
				Iter.	Ref.	Ass.	Sel.
Host system related functionality	FAU_GEN.1 Audit data generation		CC Part 2	No	No	Yes	Yes
	FAU_GEN.2 User identity association		CC Part 2	No	No	No	No
	FAU_SAR.1 Audit review		CC Part 2	No	No	Yes	No
	FAU_SAR.2 Restricted audit review		CC Part 2	No	No	No	No
	FAU_SEL.1 Selective audit		CC Part 2	No	No	Yes	Yes
	FAU_STG.1 Protected audit trail storage		CC Part 2	No	No	No	Yes
	FAU_STG.3 Action in case of possible audit data loss		CC Part 2	No	No	Yes	No
	FAU_STG.4 Prevention of audit data loss		CC Part 2	No	Yes	Yes	Yes
	FCS_CKM.1(SYM) Cryptographic key generation	FCS_CKM.1	CC Part 2	Yes	Yes	Yes	No
	FCS_CKM.1(RSA) Cryptographic key generation	FCS_CKM.1	CC Part 2	Yes	Yes	Yes	No
	FCS_CKM.1(DSA) Cryptographic key generation	FCS_CKM.1	CC Part 2	Yes	Yes	Yes	No
	FCS_CKM.2(NET) Cryptographic key distribution	FCS_CKM.2	CC Part 2	No	Yes	Yes	No
	FCS_CKM.4 Cryptographic key destruction		CC Part 2	No	No	Yes	No

Security functional group	Security functional requirement	Base security functional component	Source	Operations			
				Iter.	Ref.	Ass.	Sel.
	FCS_COP.1(NET) Cryptographic operation	FCS_COP.1	CC Part 2	Yes	Yes	Yes	No
	FCS_RNG.1(SSH-DFLT) Random number generation (Class DRG.2)	FCS_RNG.1	ECD	No	Yes	Yes	Yes
	FCS_RNG.1(SSH-FIPS) Random number generation (Class DRG.2)	FCS_RNG.1	ECD	No	Yes	Yes	Yes
	FDP_ACC.1(PSO) Subset access control	FDP_ACC.1	CC Part 2	No	No	Yes	No
	FDP_ACF.1(PSO) Security attribute based access control	FDP_ACF.1	CC Part 2	Yes	No	Yes	No
	FDP_IFC.2(NI) Complete information flow control	FDP_IFC.2	CC Part 2	Yes	No	Yes	No
	FDP_IFF.1(NI-IPTables) Simple security attributes	FDP_IFF.1	CC Part 2	Yes	No	Yes	Yes
	FDP_IFF.1(NI-ebtables) Simple security attributes	FDP_IFF.1	CC Part 2	Yes	No	Yes	Yes
	FDP_ITC.2(BA) Import of user data with security attributes	FDP_ITC.2	CC Part 2	Yes	No	Yes	No
	FDP_RIP.2 Full residual information protection		CC Part 2	No	No	No	Yes
	FDP_RIP.3 Full residual information protection of resources		ECD	No	No	No	Yes
	FIA_AFL.1 Authentication failure handling		CC Part 2	No	Yes	Yes	Yes
	FIA_ATD.1(HU) User attribute definition	FIA_ATD.1	CC Part 2	Yes	Yes	Yes	No
	FIA_ATD.1(TU) User attribute definition	FIA_ATD.1	CC Part 2	Yes	Yes	Yes	No
	FIA_SOS.1 Verification of secrets		CC Part 2	No	No	Yes	No
	FIA_UAU.1 Timing of authentication		CC Part 2	No	No	Yes	No
	FIA_UAU.5 Multiple authentication mechanisms		CC Part 2	No	No	Yes	No
	FIA_UAU.7 Protected authentication feedback		CC Part 2	No	No	Yes	No
	FIA_UID.1 Timing of identification		CC Part 2	No	No	Yes	No

Security functional group	Security functional requirement	Base security functional component	Source	Operations			
				Iter.	Ref.	Ass.	Sel.
	FIA_USB.2 Enhanced user-subject binding		ECD	No	No	Yes	No
	FPT_STM.1 Reliable time stamps		CC Part 2	No	No	No	No
	FPT_TDC.1(BA) Inter-TSF basic TSF data consistency	FPT_TDC.1	CC Part 2	Yes	No	Yes	No
	FTA_SSL.1 TSF-initiated session locking		CC Part 2	No	Yes	Yes	No
	FTA_SSL.2 User-initiated locking		CC Part 2	No	Yes	Yes	No
	FPT_ITC.1 Inter-TSF trusted channel		CC Part 2	No	Yes	Yes	Yes
Virtual machine related functionality	FDP_ACC.2(VIRT) Complete access control	FDP_ACC.2	CC Part 2	No	No	Yes	No
	FDP_ACF.1(VIRT) Security attribute based access control	FDP_ACF.1	CC Part 2	Yes	No	Yes	No
	FDP_ETC.2(VIRT) Export of user data with security attributes	FDP_ETC.2	CC Part 2	No	No	Yes	No
	FDP_IFC.2(VIRT) Complete information flow control	FDP_IFC.2	CC Part 2	Yes	No	Yes	No
	FDP_IFF.1(VIRT) Simple security attributes	FDP_IFF.1	CC Part 2	Yes	Yes	Yes	No
	FDP_ITC.2(VIRT) Import of user data with security attributes	FDP_ITC.2	CC Part 2	Yes	No	Yes	No
	FIA_UID.2(VIRT) User identification before any action	FIA_UID.2	CC Part 2	No	Yes	No	No
	FPT_TDC.1(VIRT) Inter-TSF basic TSF data consistency	FPT_TDC.1	CC Part 2	Yes	No	Yes	No
Management related functionality	FMT_MSA.1(PSO) Management of object security attributes	FMT_MSA.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MSA.1(VIRT-CACP) Management of security attributes	FMT_MSA.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MSA.1(VIRT-CIFCP) Management of security attributes	FMT_MSA.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MSA.3(PSO) Static attribute initialisation	FMT_MSA.3	CC Part 2	Yes	No	Yes	Yes
	FMT_MSA.3(NI) Static attribute initialisation	FMT_MSA.3	CC Part 2	Yes	No	Yes	Yes

Security functional group	Security functional requirement	Base security functional component	Source	Operations			
				Iter.	Ref.	Ass.	Sel.
	FMT_MSA.3(VIRT-CACP) Static attribute initialisation	FMT_MSA.3	CC Part 2	Yes	No	Yes	Yes
	FMT_MSA.3(VIRT-CIFCP) Static attribute initialisation	FMT_MSA.3	CC Part 2	Yes	Yes	Yes	Yes
	FMT_MSA.4(PSO) Security attribute value inheritance	FMT_MSA.4	CC Part 2	No	Yes	Yes	No
	FMT_MTD.1(AE) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(AS) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(AT) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(AF) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(NI) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(IAT) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(IAF) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(IAU) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(SSH) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(SSL) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_MTD.1(VIRT-COMP) Management of TSF data	FMT_MTD.1	CC Part 2	Yes	No	Yes	Yes
	FMT_REV.1(OBJ) Revocation	FMT_REV.1	CC Part 2	Yes	Yes	Yes	Yes
	FMT_REV.1(USR) Revocation	FMT_REV.1	CC Part 2	Yes	Yes	Yes	Yes
	FMT_SMF.1 Specification of management functions		CC Part 2	No	No	Yes	No
	FMT_SMR.1 Security roles		CC Part 2	No	No	Yes	No

Table 7: Security functional requirements for the TOE

6.1.1 Host system related functionality

6.1.1.1 Audit data generation (FAU_GEN.1)

- FAU_GEN.1.1** The TSF shall be able to generate an audit record of the following auditable events:
- a) Start-up and shutdown of the audit functions;
 - b) All auditable events for the **not specified** level of audit; and
 - c) **all modifications to the set of events being audited;**
 - d) **all user authentication attempts;**
 - e) **all denied accesses to objects for which the access control policies apply;**
 - f) **explicit modifications of access rights to objects covered by the access control policies; and**
 - g) **basic level of auditing for: FAU_SAR.1, FAU_SEL.1, FAU_STG.3, FAU_STG.4, FCS_CKM.1(SYM), FCS_CKM.1(RSA), FCS_CKM.1(DSA), FCS_CKM.2(NET), FCS_COP.1(NET), FDP_ACF.1(PSO), FIA_AFL.1, FIA_SOS.1, FIA_UAU.1, FIA_UAU.5, FIA_UID.1, FIA_USB.2, FPT_STM.1, FTA_SSL.1, FTA_SSL.2, FTP_ITC.1, FDP_ACF.1(VIRT), FIA_UID.2(VIRT), FMT_MSA.1(PSO), FMT_MSA.3(PSO), FMT_MSA.3(VIRT-CACP), FMT_MSA.3(VIRT-CIFCP), FMT_MSA.4(PSO), all iterations of FMT_MTD.1 except FMT_MTD.1(VIRT-COMP), FMT_REV.1(OBJ), FMT_REV.1(USR), FMT_SMF.1, FMT_SMR.1;**
 - h) **FMT_MSA.1(VIRT-CACP): assignment of disk devices and network devices to virtual machines.**

Application Note: *The auditing compliant with the basic level of auditing is under development for the following functions:*

- *Packet filter and packet filter management*
- *Management of all resource assignments to virtual machines*

- FAU_GEN.1.2** The TSF shall record within each audit record at least the following information:
- a) Date and time of the event, type of event, subject identity (if applicable), and outcome of the event; and
 - b) For each audit event type, based on the auditable event definitions of the functional components included in the PP/ST;
 - i. **User identity (if applicable); and**
 - ii. **no other audit relevant information.**

6.1.1.2 User identity association (FAU_GEN.2)

- FAU_GEN.2.1** For audit events resulting from actions of identified users, the TSF shall be able to associate each auditable event with the identity of the user that caused the event.

Application Note: *The TOE maintains a "Login UID", which is inherited by every new process spawned. This allows the TOE to identify the "real" originator of an event, regardless if he has changed his real and / or effective and filesystem UID e. g. using the su command or executing a setuid or setgid program.*

6.1.1.3 Audit review (FAU_SAR.1)

FAU_SAR.1.1 The TSF shall provide **the root user** with the capability to read **all audit information** from the audit records.

FAU_SAR.1.2 The TSF shall provide the audit records in a manner suitable for the user to interpret the information.

Application Note: *The audit records are stored in ASCII format and can therefore be read with a normal editor or pager. In addition, the TOE provides specific tools that support the interpretation of the audit trail.*

Application Note: *The audit trail is stored in a file that is readable to the root user only.*

6.1.1.4 Restricted audit review (FAU_SAR.2)

FAU_SAR.2.1 The TSF shall prohibit all users read access to the audit records, except those users that have been granted explicit read-access.

Application Note: *The protection of the audit records is based on the Unix permission bit settings defined by FDP_ACC.1(PSO) supported by FDP_ACF.1(PSO).*

6.1.1.5 Selective audit (FAU_SEL.1)

FAU_SEL.1.1 The TSF shall be able to select the set of events to be audited from the set of all auditable events based on the following attributes:

- a) **Type of audit event;**
- b) **Subject or user identity;**
- c) **Outcome (success or failure) of the audit event;**
- d) **Named object identity;**
- e) **Access types on a particular object;**
- f) **System call number;**
- g) **Virtual machine category label;**
- h) **Virtual machine resources category label.**

Application Note: *The TOE provides an application that allows specification of the audit rules which injects the rules into the kernel for enforcement. The Linux kernel auditing mechanism obtains all audit events and decides based on this rule set whether an event is forwarded to the audit daemon for storage.*

6.1.1.6 Protected audit trail storage (FAU_STG.1)

FAU_STG.1.1 The TSF shall protect the stored audit records in the audit trail from unauthorised deletion.

FAU_STG.1.2 The TSF shall be able to **prevent** unauthorised modifications to the audit records in the audit trail.

Application Note: *The protection of the audit records is based on the Unix permission bit settings defined by FDP_ACC.1(PSO) supported by FDP_ACF.1(PSO). The audit trail is readable and writeable to the root user only.*

6.1.1.7 Action in case of possible audit data loss (FAU_STG.3)

FAU_STG.3.1 The TSF shall **notify an authorized administrator** if the audit trail exceeds **an root-user selectable, pre-defined size limit of the audit trail or if any of the following condition is detected that may result in a loss of audit records : no other condition .**

Application Note: *The term "authorized administrator" refers to the user that is notified by the auditd daemon. This daemon can be configured to notify different users in different ways. The administrator of the system must ensure that the auditd is configured to send the notification to the intended recipient.*

Application Note: *The alarm generated by the TOE can be configured to be a syslog message or the execution of an administrator-specified application. This message or action of executing the application is generated when the audit trail capacity exceeds the limit defined in the auditd.conf file.*

Application Note: *The information of the threshold limit is done in the configuration file of the auditd daemon. This file is only writeable to the root user.*

6.1.1.8 Prevention of audit data loss (FAU_STG.4)

FAU_STG.4.1 The TSF shall *be able to* **ignore the audited events** and **perform one of the following administrator-defined actions:**

- a) Stop all processes that attempt to generate an audit record;**
- b) Switch to single user mode;**
- c) Halt the system**

if the audit trail is full.

Application Note: *The SFR lists all configuration possibilities that apply to the case when the audit trail is full (i.e. the disk is full). Even though the SFR mentions the "ignoring of audit events" separate from the other options, all options should be seen as equal where the root user can select one of these options.*

6.1.1.9 Cryptographic key generation (FCS_CKM.1(SYM))

FCS_CKM.1.1 The TSF shall generate *symmetric* cryptographic keys in accordance with a specified cryptographic key generation algorithm **capable of generating a random bit sequence** and specified cryptographic key sizes:

- a) 128 bits,**
- b) 168 bits,**
- c) 256 bits,**
- d) 192 bits**

that meet the following: **cryptographic key generation algorithm based on the random number generator FCS_RNG.1(SSH-DFLT) or FCS_RNG.1(SSH-FIPS) for use in OpenSSH applications.**

Application Note: *The keys are generated for AES and Triple-DES used within OpenSSH for encryption.*

6.1.1.10 Cryptographic key generation (FCS_CKM.1(RSA))

FCS_CKM.1.1 The TSF shall generate *RSA* cryptographic keys in accordance with a specified cryptographic key generation algorithm **defined in U.S. NIST FIPS PUB 186-3 appendix B.3** and specified cryptographic key sizes:

- a) **2048 bits,**
- b) **1024 bits,**
- c) **3072 bits**

that meet the following:

- a) **U.S. NIST FIPS PUB 186-3,**
- b) **no other standard.**

Application Note: *The TOE supports the generation of RSA keys for the OpenSSH host key as well as the OpenSSH user keys using the ssh-keygen(1) application.*

6.1.1.11 Cryptographic key generation (FCS_CKM.1(DSA))

FCS_CKM.1.1 The TSF shall generate *DSA* cryptographic keys in accordance with a specified cryptographic key generation algorithm **defined in U.S. NIST FIPS PUB 186-3 appendix B.1** and specified cryptographic key sizes:

- a) **L=1024, N=160 bits;**

that meet the following:

- a) **U.S. NIST FIPS PUB 186-3,**
- b) **no other standard**

Application Note: *The TOE supports the generation of DSA keys for the OpenSSH host key as well as the OpenSSH user keys using the ssh-keygen(1) application.*

6.1.1.12 Cryptographic key distribution (FCS_CKM.2(NET))

FCS_CKM.2.1 The TSF shall distribute cryptographic keys in accordance with a *the following* specified cryptographic key distribution method ~~{assignment: cryptographic key distribution method}~~ that meets the following:

- a) **Diffie-Hellman key agreement method defined for the SSH protocol by RFC4253;**
- b) **Public DSS, RSA host key exchange defined for the SSH protocol by RFC4253;**

Application Note: *DSS defined in RFC4253 for the host key exchange is compliant with DSA defined in FIPS 186-3.*

6.1.1.13 Cryptographic key destruction (FCS_CKM.4)

FCS_CKM.4.1 The TSF shall destroy cryptographic keys in accordance with a specified cryptographic key destruction method **of zerorization** that meets the following: **vendor-specific zeroization.**

Application Note:

The "vendor-specific zeroization" covers to the following concepts:

- *Memory objects: Overwriting the memory with zeros at the time the memory is released.*
- *Asymmetric key components stored in files: The object reuse functionality for objects defined with FDP_RIP.2 also covers this SFR.*

6.1.1.14 Cryptographic operation (FCS_COP.1(NET))

FCS_COP.1.1 The TSF shall perform **encryption, decryption, integrity verification, peer authentication** in accordance with ~~a specified~~*the following* cryptographic algorithms, cryptographic key sizes ~~[assignment: cryptographic key distribution method]~~ that meets the following ~~and applicable standards~~:

- a) SSH allowing the use of TDES in CBC mode with 168 bits key size, and HMAC-SHA1 defined by RFC 4253;**
- b) SSH allowing the use of AES in CBC mode with 128 bits and 256 bits key size, and HMAC-SHA1 defined by RFC 4253;**
- c) SSH allowing the use of AES in CBC mode with 192 bits key size, and HMAC-SHA1 defined by RFC 4253;**
- d) SSH allowing the use of AES in CTR mode with 128 bits, 192 bits and 256 bits key size, and HMAC-SHA1 defined by RFC 4253;**

Application Note: *The SSH protocol allows access to the console of the TOE as well as to the libvirt virtual machine management daemon.*

6.1.1.15 Random number generation (Class DRG.2) (FCS_RNG.1(SSH-DFLT))

FCS_RNG.1.1 The TSF shall provide a *deterministic* random number generator that implements:

- a) DRG2.1: If initialized with a random seed using **/dev/random as entropy source**, the internal state of the RNG shall **have 96 bits of entropy**.
- b) DRG2.2: The DRNG provides forward secrecy.
- c) DRG2.3: The DRNG provides backward secrecy.

FCS_RNG.1.2 The TSF shall provide random numbers that meet:

- a) DRG.2.4: The RNG initialized with a random seed **from the operating system seed source of /dev/random every time a new entity requires random bits** generates output for which **2**19** strings of length **1 bit** are mutually different with probability **2**-10**.
- b) DRG.2.5: The test suite A **and no other test suite** cannot distinguish the random numbers from output sequences of ideal RNGs.

Application Note:

The OpenSSH applications use the deterministic RNG from OpenSSL to generate random numbers. Every time the ssh client is invoked, the ssh-keygen application is used or a new SSH connection is processed by sshd, the deterministic random number generator is seeded with data from /dev/random. This SFR covers the default DRNG provided by OpenSSL.

6.1.1.16 Random number generation (Class DRG.2) (FCS_RNG.1(SSH-FIPS))

- FCS_RNG.1.1** The TSF shall provide a *deterministic* random number generator *using the ANSI X9.31 Appendix A2.4 standard with AES-128 core* that implements:
- a) DRG2.1: If initialized with a random seed using **/dev/random as entropy source**, the internal state of the RNG shall **have 96 bits of entropy**.
 - b) DRG2.2: The DRNG provides forward secrecy.
 - c) DRG2.3: The DRNG provides backward secrecy.
- FCS_RNG.1.2** The TSF shall provide random numbers that meet:
- a) DRG.2.4: The RNG initialized with a random seed **from the operating system seed source of /dev/random every time a new entity requires random bits** generates output for which **2**19** strings of length **1 bit** are mutually different with probability **2**-10**.
 - b) DRG.2.5: The test suite A **and no other test suite** cannot distinguish the random numbers from output sequences of ideal RNGs.

Application Note:

The OpenSSH applications use the deterministic RNG from OpenSSL to generate random numbers. Every time the ssh client is invoked, the ssh-keygen application is used or a new SSH connection is processed by sshd, the deterministic random number generator is seeded with data from /dev/random. This SFR covers the FIPS 140-2 compliant DRNG provided by OpenSSL.

6.1.1.17 Subset access control (FDP_ACC.1(PSO))

- FDP_ACC.1.1** The TSF shall enforce the **Persistent Storage Object Access Control Policy** on
- a) **Subjects: all subjects defined with the Security Policy Model;**
 - b) **Objects:**
 - i. **Persistent Storage Objects of the following type : file system objects;**
 - ii. **no other storage objects;**
 - c) **Operations: read, write, execute (regular files), search (directories).**

6.1.1.18 Security attribute based access control (FDP_ACF.1(PSO))

- FDP_ACF.1.1** The TSF shall enforce the **Persistent Storage Object Access Control Policy** to objects based on the following:
- a) **Subject security attributes: file system UID, file system GID, supplemental GIDs;**
 - b) **Object security attributes: owning UID, owning GID;**
 - c) **Access control security attributes maintained for each file system object governing access to that object:**
 - i. **ACL for specific UIDs (ACL_USER),**
 - ii. **ACL for specific GIDs (ACL_GROUP),**
 - iii. **Maximum ACL for the file system object (ACL_MASK),**

- iv. **Permission bits for the owning UID (equals to ACL_USER_OBJ when using ACLs),**
- v. **Permission bits for the owning GID (equals to ACL_GROUP_OBJ when using ACLs),**
- vi. **Permission bits for "world" (equals to ACL_OTHER when using ACLs),**
- vii. **The following permission bits: read, write, execute (for files), search (for directories),**
- viii. **The following access rights applicable to the file system object: SAVETXT (directories), immutable (files),**

d) Access control security attributes maintained for each partition holding a file system: read-only;

FDP_ACF.1.2

The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed:

A subject must have search permission for every element of the pathname and the requested access for the object. A subject has a specific type access to an object if one of the following rules hold (the order of the rules is applicable on a first-match basis):

- a) **The subject's filesystem UID is identical with the owning UID of the object and the requested type of access is within the permission bits defined for the owning UID (permission bits) or by ACL_USER_OBJ (ACLs); or**
- b) **ACLs: The subject's filesystem UID is identical with the UID specified with ACL_USER of the object and the requested type of access is within the permission bits defined in ACL_USER; or**
- c) **The subject's filesystem GID or one of the subject's supplemental GIDs identical with the owning GID and the requested type of access is within the permission bits defined for the owning GID (permission bits), or by ACL_GROUP_OBJ when there is no ACL_MASK entry (ACLs), or by the ACL_MASK entry (ACLs); or**
- d) **ACLs: The subject's filesystem GID or one of the subject's supplemental GIDs is identical with the GID specified with ACL_GROUP of the object and the requested type of access is within the permission bits defined in ACL_GROUP; or**
- e) **The requested type of access is within the permission bits defined for "world" (permission bits) or by ACL_OTHER (ACLs).**

Application Note: *The permission bits and the ACLs are inherently consistent as the TOE assigns the permission bits to ACLs when ACLs are used. Without any ACLs specified for an object, the TOE only uses the permission bits. If at least one ACL is present or when the ACL management tools are applied for objects even without any ACL set, the permission bits are interpreted as outlined above: the ACL entry of ACL_USER_OBJ contains the owning UID permission bits, the ACL entry of ACL_GROUP_OBJ contains the owning GID permission bits, and the ACL entry of ACL_OTHER contains the permission bits for "world". The ACL entries of ACL_USER_OBJ, ACL_GROUP_OBJ and ACL_OTHER are only a different representation of the permission bits to users, they are not separate attributes in addition to permission bits. The explicit specification of ACL_USER_OBJ, ACL_GROUP_OBJ and ACL_OTHER in the rule set above in addition to the permission bits is only intended to aid the evaluator or reader in understanding the overall ruleset.*

Application Note: *Due to the fact that the permission bits are an inherent part of the ACLs, there is no precedence issue between permission bits and ACLs.*

FDP_ACF.1.3 The TSF shall explicitly authorise access of subjects to objects based on the following additional rules: **all operations except execute operations are allowed for the subject with the file system UID of zero - the execute permission is granted if the file system object object is marked with at least one executable bit in its permission settings.**

FDP_ACF.1.4 The TSF shall explicitly deny access of subjects to named objects based on the following rules:

- a) **Any file system object in a file system that is mounted as read-only cannot be modified, created or removed,**
- b) **Any file system object marked as immutable cannot be modified or removed,**
- c) **Any file system object stored in a directory marked with the SAVETXT bit cannot be modified or removed by subjects whose file system UID is not equal to the owning UID of the file system object.**

6.1.1.19 Complete information flow control (FDP_IFC.2(NI))

FDP_IFC.2.1 The TSF shall enforce the **Network Information Flow Control Policy** on

- a) **Subjects:**
 - i. **unauthenticated external IT entities that send and receive information mediated by the TOE;**
 - ii. **standard Linux processes and processes providing a virtual machine environment that send and receive information mediated by the TOE;**
- b) **Information:**
 - i. **Network data routed through the TOE;**
 - ii. **No other information;**

and all operations that cause that information to flow to and from subjects covered by the SFP.

FDP_IFC.2.2 The TSF shall ensure that all operations that cause any information in the TOE to flow to and from any subject in the TOE are covered by an information flow control SFP.

Application Note: *The SFR applies only to host systems which implements the packet filtering functionality.*

6.1.1.20 Simple security attributes (FDP_IFF.1(NI-IPTables))

FDP_IFF.1.1 The TSF shall enforce the **Network Information Flow Control Policy** based on the following types of subject and information security attributes:

- a) **Information security attribute: the logical or physical network interface through which the network data entered the TOE;**
- b) **TCP/IP information security attributes:**
 - i. **Source and destination IP address,**

- ii. **Source and destination TCP port number,**
 - iii. **Source and destination UDP port number,**
 - iv. **Network protocol of IP, TCP, UDP, ICMP**
 - v. **TCP header flags of SYN, ACK, FIN, RST, URG, PSH**
 - vi. **TCP sequence numbers for IPv4;**
- c) **Information security attribute: IP packet identified as to be routed by the TSF.**

Application Note: *Stateful packet filtering is only implemented for IPv4.*

FDP_IFF.1.2 The TSF shall permit an information flow between a controlled subject and controlled information via a controlled operation if the following rules hold: **If the packet filter matches the analyzed packet and the rule accepts the packet, the packet is forwarded according to the network protocol stack's behavior .**

FDP_IFF.1.3 The TSF shall enforce the **following rules:**
Identification of network data using one or more of the following concepts:

- a) **Information security attribute matching;**
- b) **Matching based on the state of a TCP connection, Statistical analysis matching;**

Performing one or more of the following actions with identified network data:

- a) **Discard the network data without any further processing, with sending a notification to the sender;**
- b) **Allow the network data to be processed unaltered by the TOE according to the routing information maintained by the TOE;**
- c) **No other actions.**

FDP_IFF.1.4 The TSF shall explicitly authorise an information flow based on the following rules: **If the network data is not matched by the rule set and the default rule of the packet filter is ACCEPT then the data is forwarded unaltered based on the normal operation of the host system's networking stack .**

FDP_IFF.1.5 The TSF shall explicitly deny an information flow based on the following rules: **If the network data is not matched by the rule set, one of the following default rules applies:**

- a) **DROP: the data is discarded;**
- b) **REJECT: then the data is discarded and a notification is returned to the sender.**

Application Note: *The default rule is configurable where exactly one of the above mentioned default rules can be selected at any given time.*

6.1.1.21 Simple security attributes (FDP_IFF.1(NI-ebtables))

FDP_IFF.1.1 The TSF shall enforce the **Network Information Flow Control Policy** based on the following types of subject and information security attributes:

- a) **Information security attribute: the logical or physical network interface through which the network data entered the TOE;**

b) TCP/IP information security attributes:

- i. Source and destination IP address,**
- ii. Source and destination TCP port number,**
- iii. Source and destination UDP port number,**
- iv. Network protocol of IP, TCP, UDP**
- v. TCP header flags of TOS.**

- FDP_IFF.1.2** The TSF shall permit an information flow between a controlled subject and controlled information via a controlled operation if the following rules hold: **If the packet filter matches the analyzed packet and the rule accepts the packet, the packet is forwarded according to the network protocol stack's behavior .**
- FDP_IFF.1.3** The TSF shall enforce the **following rules:**
Identification of network data using one or more of the following concepts:
- a) Information security attribute matching;**
 - b) No other concepts;**
- Performing one or more of the following actions with identified network data:**
- a) Discard the network data without any further processing;**
 - b) Allow the network data to be processed unaltered by the TOE according to the routing information maintained by the TOE;**
 - c) No other actions.**
- FDP_IFF.1.4** The TSF shall explicitly authorise an information flow based on the following rules: **If the network data is not matched by the rule set and the default rule of the packet filter is ACCEPT then the data is forwarded unaltered based on the normal operation of the host system's networking stack .**
- FDP_IFF.1.5** The TSF shall explicitly deny an information flow based on the following rules: **If the network data is not matched by the rule set, and the default rule of the packet filter is DROP then the data is discarded.**

Application Note: *The default rule is configurable where exactly one of the above mentioned default rules can be selected at any given time.*

6.1.1.22 Import of user data with security attributes (FDP_ITC.2(BA))

- FDP_ITC.2.1** The TSF shall enforce the **Network Information Flow Control Policy** when importing user data, controlled under the SFP, from outside of the TOE.
- FDP_ITC.2.2** The TSF shall use the security attributes associated with the imported user data.
- FDP_ITC.2.3** The TSF shall ensure that the protocol used provides for the unambiguous association between the security attributes and the user data received.
- FDP_ITC.2.4** The TSF shall ensure that interpretation of the security attributes of the imported user data is as intended by the source of the user data.
- FDP_ITC.2.5** The TSF shall enforce the following rules when importing user data controlled under the SFP from outside the TOE: **No additional importation control rules.**

6.1.1.23 Full residual information protection (FDP_RIP.2)

FDP_RIP.2.1 The TSF shall ensure that any previous information content of a resource is made unavailable upon the **allocation of the resource to** all objects.

6.1.1.24 Full residual information protection of resources (FDP_RIP.3)

FDP_RIP.3.1 The TSF shall ensure that any previous information content of a resource is made unavailable upon the **allocation of the resource to** all subjects or users.

Application Note: *The subject is represented by the data structures inside the kernel forming a process: all data structures anchored in the task_struct. The user is represented by its attributes defined by FIA_ATD.1.*

6.1.1.25 Authentication failure handling (FIA_AFL.1)

FIA_AFL.1.1 The TSF shall detect when **an administrator-configurable number of** unsuccessful authentication attempts *for the authentication method of password-based authentication* occur related to **consecutive unsuccessful authentication attempts**.

FIA_AFL.1.2 When the defined number of unsuccessful authentication attempts has been **met, surpassed**, the TSF shall

- a) **For all administrator accounts, "disable" the account for an authorized administrator configurable time period such that there can be no more than ten attempts per minute.**
- b) **For all other accounts, disable the user logon account until it is re-enabled by the authorized administrator.**
- c) **For all disabled accounts, any response to an authentication attempt given to the user shall not be based on the result of that authentication attempt.**

6.1.1.26 User attribute definition (FIA_ATD.1(HU))

FIA_ATD.1.1 The TSF shall maintain the following list of security attributes belonging to individual *human* users:

- a) **User identifier;**
- b) **Group memberships;**
- c) **User password;**
- d) **Software token verification data;**
- e) **Security roles;**
- f) **no other security attributes.**

Application Note: *Please see the application note for FIA_UAU.5 for a list of token-based authentication mechanisms and their associated tokens.*

6.1.1.27 User attribute definition (FIA_ATD.1(TU))

FIA_ATD.1.1 The TSF shall maintain the following list of security attributes belonging to individual *technical* users:

- a) **the logical or physical network interface through which the network data entered the TOE;**
- b) **identity of the logical or physical external interface through which the user connected to the TOE;**
- c) **no other security attributes.**

Application Note: *Bullet a) of this SFR relates to FDP_IFC.2(NI) and the supporting information flow control rules specified with the iterations of FDP_IFF.1. In the Common Criteria scheme, external entities are always considered to be users. Therefore, every network data entity must be specified as user in this PP.*

6.1.1.28 Verification of secrets (FIA_SOS.1)

FIA_SOS.1.1 The TSF shall provide a mechanism to verify that secrets meet **the following quality metric: the probability that a secret can be obtained by an attacker during the lifetime of the secret is less than 2^{-20} .**

Application Note: *The TOE password change is implemented using the PAM library. The PAM module `pam_passwordqc.so` allows the specification of the quality of new passwords. The evaluated configuration requires a configuration of the PAM-based password change mechanism that meets the above mentioned criteria.*

Application Note: *The Evaluated Configuration Guide contains configuration suggestions for the password quality mechanism that covers the above mentioned probability. These configuration suggestions assume the worst-case scenario when attacking these settings.*

Application Note: *For key-based authentication methods, the evaluation of the RSA and DSA keys used for the SSH protocol will show the maximum lifetime of a key depending on its size.*

6.1.1.29 Timing of authentication (FIA_UAU.1)

FIA_UAU.1.1 The TSF shall allow

- a) **the information flow covered by the Network Information Flow Control Policy;**
- b) **Local console log-in: banner information;**
- c) **SSH log-in: obtaining the list of allowed authentication methods;**

on behalf of the user to be performed before the user is authenticated.

FIA_UAU.1.2 The TSF shall require each user to be successfully authenticated before allowing any other TSF-mediated actions on behalf of that user.

6.1.1.30 Multiple authentication mechanisms (FIA_UAU.5)

FIA_UAU.5.1 The TSF shall provide **the following authentication mechanisms:**

- a) **Authentication based on username and password;**
- b) **Authentication based on software token verification data;**
- c) **No other authentication mechanisms**

to support user authentication.

Application Note: *The TOE is able to maintain the following types of software tokens and their verification data:*

- *SSH user keys: The TOE as server part is able to store the public part of the SSH user key for the user account the user wants to access.*

FIA_UAU.5.2 The TSF shall authenticate any user's claimed identity according to the **following rules:**

- a) Authentication based on username and password is performed for TOE-originated requests and credentials stored by the TSF;**
- b) Authentication based on software token verification data is performed for TOE-originated requests;**
- c) For SSH, both, the password-based and key-based authentication methods can be enabled at the same time. In this case, the key-based authentication method is tried before the password-based authentication. If the key-based authentication succeeds, the user is authenticated. If the key-based authentication fails, the password-based authentication is applied. If the password-based authentication fails, the user login request is denied.**

6.1.1.31 Protected authentication feedback (FIA_UAU.7)

FIA_UAU.7.1 The TSF shall provide only **obscured feedback** to the user while the authentication is in progress.

6.1.1.32 Timing of identification (FIA_UID.1)

FIA_UID.1.1 The TSF shall allow

- a) Console log-in: banner information;**
- b) SSH log-in: obtaining the list of allowed authentication methods;**

on behalf of the user to be performed before the user is identified.

FIA_UID.1.2 The TSF shall require each user to be successfully identified before allowing any other TSF-mediated actions on behalf of that user.

6.1.1.33 Enhanced user-subject binding (FIA_USB.2)

FIA_USB.2.1 The TSF shall associate the following user security attributes with subjects acting on the behalf of that user:

- a) The user identity that is associated with auditable events;**
- b) The user security attributes that are used to enforce the Persistent Storage Object Access Control Policy;**
- c) The software token that can be used for subsequent identification and authentication with the TSF or other remote IT systems;**
- d) Active roles;**
- e) Active groups;**

f) The user security attributes that are used to enforce the Compartment Access Control Policy.

FIA_USB.2.2

The TSF shall enforce the following rules on the initial association of user security attributes with subjects acting on the behalf of users:

- a) Upon successful identification and authentication, the login UID, the real UID, the filesystem UID and the effective UID shall be those specified in the user entry for the user that has authenticated successfully;**
- b) Upon successful identification and authentication, the real GID, the filesystem GID and the effective GID shall be those specified via the primary group membership attribute in the user entry;**
- c) Upon successful identification and authentication, the supplemental GIDs shall be those specified via the supplemental group membership assignment for the user entry;**
- d) Upon instantiating a virtual machine, the category selected by the virtual machine management daemon is associated with the process representing a virtual machine environment.**

Application Note: *The various subject UIDs are all derived from the same numeric UID per user entry stored in the /etc/passwd file.*

Application Note: *The various subject GIDs except the supplemental GIDs are all derived from the same numeric GID per user entry stored in the /etc/passwd file.*

Application Note: *The subject's supplemental GIDs are derived from the username to group name mappings in the /etc/group file. As the TOE only maintains numeric IDs for subjects, the username and the group names need to be converted before instantiating the subject. The username to UID mapping is provided in /etc/passwd and the group name to GID mapping is provided in /etc/group.*

Application Note: *The libvirt virtual machine management daemon automatically identifies a yet unused category during the initial configuration of a virtual machine and stores the information together with the virtual machine configuration in /etc/libvirt/qemu/. This category is applied every time the virtual machine is instantiated.*

FIA_USB.2.3

The TSF shall enforce the following rules governing changes to the user security attributes associated with subjects acting on the behalf of users:

- a) The effective and filesystem UID of a subject can be changed by the use of an executable with the SETUID bit set. In this case the program is executed with the effective and filesystem UID of the owning UID of the file storing the program. These newly set effective and filesystem UIDs are used for the DAC permission validation. The real and login UID remain unchanged.**
- b) The effective and filesystem GID of a subject can be changed by the use of an executable with the SETGID bit set. In this case the program is executed with the effective and filesystem GID of the owning GID of the file storing the program. These newly set effective and filesystem GIDs are used for the DAC permission validation. The real GID remains unchanged.**
- c) The real, effective and filesystem UID of a subject can be changed by the use of the set*uid system call family for the calling application. These system calls are restricted to the root user.**

- d) **The real, effective and filesystem GID of a subject can be changed by the use of the set*gid system call family for the calling application. These system calls are restricted to the root user.**
- e) **The set of supplemental GIDs of a subject can be changed by the use of the setgroups system call for the calling application. This system call is restricted to the root user.**

Application Note: *The applications "su" allows the calling user to change the filesystem and effective UID either to root or to other users provided the authentication to "su" was successful. The su application uses the SETUID bit with the owning UID of root as well as the set*uid system calls to change to other UIDs before spawning a new shell. As the su application rests on the above mentioned mechanisms, it is not listed as a separate mechanism to modify the calling user's UIDs.*

Application Note: *The login UID is set by the PAM modules by inserting the intended UID into the /proc/<PID>/loginuid file. This file can be written to only by subjects executing with the effective UID of zero (root) and only for the calling process' own loginuid file. However, there is no application except the PAM modules which access that proc file which implies that the login UID remains unchanged after login when operating the TOE. Authorized administrators are not intended to access that proc file.*

Application Note: *The category applicable to a virtual machine cannot be modified.*

FIA_USB.2.4 The TSF shall enforce the following rules for the assignment of subject security attributes not derived from user security attributes when a subject is created:
No rules.

6.1.1.34 Reliable time stamps (FPT_STM.1)

FPT_STM.1.1 The TSF shall be able to provide reliable time stamps.

6.1.1.35 Inter-TSF basic TSF data consistency (FPT_TDC.1(BA))

FPT_TDC.1.1 The TSF shall provide the capability to consistently interpret **the following data types:**

- a) **Packet filter: protocol headers for the network protocols covered by the packet filter;**

when shared between the TSF and another trusted IT product.

FPT_TDC.1.2 The TSF shall use **the following interpretation rules:**

- a) **Packet filter: protocol headers specification provided in RFC 791 (IP), RFC 793 (TCP), RFC 768 (UDP), RFC 792 (ICMP);**

when interpreting the TSF data from another trusted IT product.

6.1.1.36 TSF-initiated session locking (FTA_SSL.1)

FTA_SSL.1.1 The TSF shall lock an interactive session *to a human user maintained by the TSF* after **an administrator-configurable time interval of user inactivity** by:

- a) clearing or overwriting *TSF controlled* display devices, making the current contents unreadable;
- b) disabling any activity of the user's *TSF controlled* data access/ *TSF controlled* display devices other than unlocking the session.

Application Note: *The management aspect of configuring the time interval is covered by FMT_MTD.1(SSL).*

- FTA_SSL.1.2** The TSF shall require the following events to occur prior to unlocking the session:
- a) **Successful re-authentication with the credentials of the user owning the session using password based authentication;**
 - b) **No other events .**

6.1.1.37 User-initiated locking (FTA_SSL.2)

- FTA_SSL.2.1** The TSF shall lock an interactive session *maintained by the TSF*, by:
- a) clearing or overwriting *TSF controlled* display devices, making the current contents unreadable;
 - b) disabling any activity of the user's *TSF controlled* data access/ *TSF controlled* display devices other than unlocking the session.

- FTA_SSL.2.2** The TSF shall require the following events to occur prior to unlocking the session:
- a) **Successful re-authentication with the credentials of the user owning the session using password based authentication;**
 - b) **No other events .**

6.1.1.38 Inter-TSF trusted channel (FTP_ITC.1)

- FTP_ITC.1.1** The TSF shall provide a communication channel between itself and another trusted IT product that is logically distinct from other communication channels and provides assured identification of its end points and protection of the channel data from modification or *and disclosure using the following mechanisms:*
- a) *Cryptographically-protected communication channel using SSH;*
 - b) *No other trusted communication channel*

- FTP_ITC.1.2** The TSF shall permit **another trusted IT product** to initiate communication via the trusted channel.

- FTP_ITC.1.3** The TSF shall initiate communication via the trusted channel for **all security functions specified in the ST that interact with remote trusted IT systems and no other conditions or functions.**

6.1.2 Virtual machine related functionality

6.1.2.1 Complete access control (FDP_ACC.2(VIRT))

- FDP_ACC.2.1** The TSF shall enforce the **Compartment Access Control Policy** on
- a) **Subjects: compartments;**
 - b) **Objects: all virtual machine resources defined with the Security Policy Model**
- and all operations among subjects and objects covered by the SFP.

Application Note: *Compartments are implemented as processes executing concurrently with standard Linux processes. These compartment processes provide the virtual machine environment.*

FDP_ACC.2.2 The TSF shall ensure that all operations between any subject controlled by the TSF and any object controlled by the TSF are covered by an access control SFP.

6.1.2.2 Security attribute based access control (FDP_ACF.1(VIRT))

FDP_ACF.1.1 The TSF shall enforce the **Compartment Access Control Policy** to objects based on the following:

a) Subject security attributes:

- i. SELinux-based: Non-hierarchical category part of the SELinux label;**
- ii. DAC-file-system-based: user ID of the QEMU virtualization process;**

b) Object security attributes:

- i. SELinux-based: Disks, network devices: non-hierarchical category part of the SELinux label;**
- ii. DAC-file-system-based: ownership and DAC permissions of file system object;**

c) No other security attribute.

FDP_ACF.1.2 The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed:

- a) SELinux-based: Access of a compartment to an object is allowed when the requested mode of access is allowed for the compartment by the compartment access control permission settings for that object;**
- b) DAC-file-system-based: Access control rules defined with FDP_ACF.1(PSO);**

Application Note: *The TOE provides the following types of access control:*

- *SELinux-based: If the category of the access-requesting virtual machine matches the category of the accessed resource, access is granted. Otherwise, access is denied to the resource.*
- *DAC: Regular Unix permission bits and ACLs as outlined in the preceding SFRs.*

FDP_ACF.1.3 The TSF shall explicitly authorise access of subjects to objects based on the following additional rules: **No additional rules.**

Application Note: *As the virtual machine processes never execute with root privileges, the DAC override rule for root does not apply here.*

FDP_ACF.1.4 The TSF shall explicitly deny access of subjects to objects based on the following additional rules: **No additional rules.**

6.1.2.3 Export of user data with security attributes (FDP_ETC.2(VIRT))

FDP_ETC.2.1 The TSF shall enforce the **Compartment Access Control Policy** and **Compartment Information Flow Control Policy** when exporting user data, controlled under the SFP(s), outside of the TOE.

FDP_ETC.2.2 The TSF shall export the user data with the user data's associated security attributes.

FDP_ETC.2.3 The TSF shall ensure that the security attributes, when exported outside the TOE, are unambiguously associated with the exported user data.

FDP_ETC.2.4 The TSF shall enforce the following rules when user data is exported from the TOE: **The host system ensures that the source IP-address is equal to the IP-address assigned to the virtual machine initiating the data export using the network.**

6.1.2.4 Complete information flow control (FDP_IFC.2(VIRT))

FDP_IFC.2.1 The TSF shall enforce the **Compartment Information Flow Control Policy** on

a) Subjects:

- i. Compartments;**
- ii. External entities;**
- iii. No other entities;**

b) Information:

- i. User data belonging to compartments;**
- ii. User data belonging to subjects outside of compartments;**
- iii. TSF data;**
- iv. No additional information**

and all operations that cause that information to flow to and from subjects covered by the SFP.

FDP_IFC.2.2 The TSF shall ensure that all operations that cause any information in the TOE to flow to and from any subject in the TOE are covered by an information flow control SFP.

6.1.2.5 Simple security attributes (FDP_IFF.1(VIRT))

FDP_IFF.1.1 The TSF shall enforce the **Compartment Information Flow Control Policy** based on the following types of subject and information security attributes:

a) Subject security attributes:

- i. Process ID of the process providing a virtual machine environment;**
- ii. No additional attributes;**

b) Information security attributes:

- i. Any attribute indicating the originating process providing a virtual machine environment;**
- ii. No TSF data security attributes;**
- iii. No additional information security attributes.**

FDP_IFF.1.2 The TSF shall permit an information flow between a controlled subject and controlled information via a controlled operation if the following rules hold: **No information flow permitted.**

Application Note: *Although virtual machines are implemented using Linux processes, the virtualization mechanism does not allow the use of any Linux inter-process communication by such a process. Therefore, a process cannot communicate with other virtual machines or other Linux processes except via the network channel to the guest system and the VNC network channel to the console of the virtual machine environment. These network channels are mediated by the TOE.*

- FDP_IFF.1.3** The TSF shall enforce the **no additional information flow control SFP rules**.
- FDP_IFF.1.4** The TSF shall explicitly authorize an information flow based on the following rules:
None.
- FDP_IFF.1.5** The TSF shall explicitly deny an information flow based on the following rules:
None.

6.1.2.6 Import of user data with security attributes (FDP_ITC.2(VIRT))

- FDP_ITC.2.1** The TSF shall enforce the **Compartment Access Control Policy and Compartment Information Flow Control Policy** when importing user data, controlled under the SFP, from outside of the TOE.
- FDP_ITC.2.2** The TSF shall use the security attributes associated with the imported user data.
- FDP_ITC.2.3** The TSF shall ensure that the protocol used provides for the unambiguous association between the security attributes and the user data received.
- FDP_ITC.2.4** The TSF shall ensure that interpretation of the security attributes of the imported user data is as intended by the source of the user data.
- FDP_ITC.2.5** The TSF shall enforce the following rules when importing user data controlled under the SFP from outside the TOE: **The host system ensures that the destination IP-address of a network packet is equal to the IP-address assigned to the virtual machine receiving the data using the network.**

6.1.2.7 User identification before any action (FIA_UID.2(VIRT))

- FIA_UID.2.1** The TSF shall require each *compartment* user to be successfully identified before allowing any other TSF-mediated actions on behalf of that user.

Application Note: *This SFR applies to the process ID of the processes providing a virtual machine environment.*

6.1.2.8 Inter-TSF basic TSF data consistency (FPT_TDC.1(VIRT))

- FPT_TDC.1.1** The TSF shall provide the capability to consistently interpret **access control and information flow control-related security attributes** when shared between the TSF and another trusted IT product.
- FPT_TDC.1.2** The TSF shall use **the IP addresses part of the network packet transmitted by the TSF as specified in RFC 791** when interpreting the TSF data from another trusted IT product.

6.1.3 Management related functionality

6.1.3.1 Management of object security attributes (FMT_MSA.1(PSO))

FMT_MSA.1.1 The TSF shall enforce the **Persistent Storage Object Access Control Policy** to restrict the ability to **modify, change_default, query** the security attributes **of the objects covered by the SFP** to **the owner of the object and users with processes granted the CAP_DAC_OVERRIDE capability**.

6.1.3.2 Management of security attributes (FMT_MSA.1(VIRT-CACP))

FMT_MSA.1.1 The TSF shall enforce the **Compartment Access Control Policy** to restrict the ability to **change_default, query, modify, delete** the security attributes **of subjects and objects covered by the SFP** to **the authorized administrator of the virtual machine configuration**.

Application Note: *The authorized administrator is the user allowed to access the libvirt daemon and configure virtual machine parameters.*

6.1.3.3 Management of security attributes (FMT_MSA.1(VIRT-CIFCP))

FMT_MSA.1.1 The TSF shall enforce the **Compartment Information Flow Control Policy** to restrict the ability to **change_default, query, modify, delete** the security attributes **of subjects and information covered by the SFP** to **the authorized administrator of the virtual machine configuration**.

Application Note: *The authorized administrator is the user allowed to access the libvirt daemon and configure virtual machine parameters.*

6.1.3.4 Static attribute initialisation (FMT_MSA.3(PSO))

FMT_MSA.3.1 The TSF shall enforce the **Persistent Storage Object Access Control Policy** to provide **restrictive** default values for security attributes that are used to enforce the SFP.

FMT_MSA.3.2 The TSF shall allow the

- a) root user for a global setting applied during logon;**
- b) each user for a setting applicable to his processes**
- c) users with write permissions to a directory for setting default ACLs**

to specify alternative initial values to override the default values when an object or information is created.

Application Note: *The global default value for permission bits is specified with the umask value which specifies the permission bits for newly created objects. This value has an initial setting of 022 or the value specified in /etc/login.defs. Only the root user can manage that initial value as this file is writeable to root only. Users can change their umask value at any time using the umask(2) system call. For ACLs, the default ACL is provided for for the root directory which, in case of absence of a default ACL entry is consistent with the umask.*

6.1.3.5 Static attribute initialisation (FMT_MSA.3(NI))

FMT_MSA.3.1 The TSF shall enforce the **Network Information Flow Control Policy** to provide **permissive** default values for security attributes that are used to enforce the SFP.

FMT_MSA.3.2 The TSF shall allow the **users with processes granted the CAP_NET_ADMIN capability** to specify alternative initial values to override the default values when an object or information is created.

Application Note: *The default value specified in this SFR applies to the default for the packet filter after boot. The administrator can configure alternative default values as outlined in FDP_IFF.1(NI-IPTables) as well as FDP_IFF.1(NI-eatables).*

Application Note: *The iptables and eatables commands use a netlink interfact to the kernel which requires that the caller possesses the CAP_NET_ADMIN capability.*

6.1.3.6 Static attribute initialisation (FMT_MSA.3(VIRT-CACP))

FMT_MSA.3.1 The TSF shall enforce the **Compartment Access Control Policy** to provide **restrictive** default values for security attributes that are used to enforce the SFP.

FMT_MSA.3.2 The TSF shall allow the

- a) **root user to alter the SELinux policy file;**
- b) **users with processes granted the CAP_SYS_ADMIN capability to load the SELinux policy into the kernel and alter the object SELinux label;**
- c) **users with processes granted the CAP_MAC_ADMIN capability to altering subject SELinux label.**

to specify alternative initial values to override the default values when an object or information is created.

Application Note: *The SFR indicates that administrative users of the TOE are technically able to change the rules for this policy enforced with an SELinux policy. Nevertheless, in the evaluated configuration, the administrative user shall not modify this SELinux policy.*

6.1.3.7 Static attribute initialisation (FMT_MSA.3(VIRT-CIFCP))

FMT_MSA.3.1 The TSF shall enforce the **Compartment Information Flow Control Policy** to provide **restrictive** default values for security attributes that are used to enforce the SFP.

FMT_MSA.3.2 The TSF shall allow the **nobody** to specify alternative initial values to override the default values when an object or information is created.

6.1.3.8 Security attribute value inheritance (FMT_MSA.4(PSO))

FMT_MSA.4.1 The TSF shall use the following rules to set the value of security attributes *for Persistent Storage Objects*:

- a) **The newly created object's owning UID is set to the effective UID of the calling subject;**

- b) **The newly created object's owning GID is set to the effective GID of the calling subject with the following exception for file system objects: if the parent directory holding the newly created file system object is marked with the SETGID permission bit, the owning GID of the newly created file system object is set to the owning GID of the parent directory;**
- c) **The newly created object's permission bits are derived from the calling subject's umask value by masking out the umask bits from the permission bit set granting full access;**
- d) **The newly created object's ACLs are derived from the default ACL specified for the parent directory the newly created file system object is stored in, if existant. Otherwise, no ACL is set.**
- e) **Permission masks (for key retention objects) are set to allow all operations for the possessor except the set attribute permission. No other permission is set for the permission masks.**

6.1.3.9 Management of TSF data (FMT_MTD.1(AE))

FMT_MTD.1.1 The TSF shall restrict the ability to **query, modify** the **set of audited events to processes with the capability CAP_AUDIT_CONTROL**.

Application Note: *This SFR applies to FAU_SEL.1.*

Application Note: *Using the audit tools which in turn use the netlink interface, an administrator can configure the audit rules.*

6.1.3.10 Management of TSF data (FMT_MTD.1(AS))

FMT_MTD.1.1 The TSF shall restrict the ability to **clear, delete, configure the storage location** the **audit storage to the root user**.

Application Note: *This SFR applies to FAU_STG.1 where the directory used for storing the audit trail is configured.*

Application Note: *The configuration of these parameters is performed with the configuration file /etc/auditd/auditd.conf which is writeable to the root user only.*

6.1.3.11 Management of TSF data (FMT_MTD.1(AT))

FMT_MTD.1.1 The TSF shall restrict the ability to **modify** the

- a) **threshold of the audit trail when an action is performed;**
- b) **action when the threshold is reached**

to **the root user**.

Application Note: *This SFR applies to FAU_STG.3.*

Application Note: *The configuration of these parameters is performed with the configuration file /etc/auditd/auditd.conf which is writeable to the root user only.*

6.1.3.12 Management of TSF data (FMT_MTD.1(AF))

FMT_MTD.1.1 The TSF shall restrict the ability to **modify the actions to be taken in case of audit storage failure to the root user.**

Application Note: *This SFR applies to FAU_STG.4.*

Application Note: *The configuration of these parameters is performed with the configuration file /etc/auditd/auditd.conf which is writeable to the root user only.*

6.1.3.13 Management of TSF data (FMT_MTD.1(NI))

FMT_MTD.1.1 The TSF shall restrict the ability to **query, modify, delete, change_default** the **security attributes for the rules governing the**

- a) identification of network packets by the packet filter;**
- b) actions performed on the identified network packets by the packet filter;**

to **users with processes granted the CAP_NET_ADMIN capability.**

Application Note: *This SFR applies to FDP_IFF.1(NI).*

Application Note: *The iptables and ebtables commands use a netlink interfact to the kernel which requires that the caller possesses the CAP_NET_ADMIN capability.*

6.1.3.14 Management of TSF data (FMT_MTD.1(IAT))

FMT_MTD.1.1 The TSF shall restrict the ability to **modify the threshold for unsuccessful authentication attempts to the root user.**

Application Note: *This SFR applies to FIA_AFL.1.*

Application Note: *The configuration of these parameters is performed with the PAM configuration files which are writeable to the root user only.*

6.1.3.15 Management of TSF data (FMT_MTD.1(IAF))

FMT_MTD.1.1 The TSF shall restrict the ability to **re-enable the authentication to the account subject to authentication failure to the root user.**

Application Note: *This SFR applies to FIA_AFL.1.*

Application Note: *The account locking information is stored in the file /var/log/tallylog. Using the pam_tally2 application which modifies this file, the account can be unlocked. The DAC permissions of that file ensure that only the root user can write to it.*

6.1.3.16 Management of TSF data (FMT_MTD.1(IAU))

FMT_MTD.1.1 The TSF shall restrict the ability to **modify, delete, initialize the user security attributes to**

- a) the root user,**
- b) users authorized to modify their own authentication data**

Application Note: *This SFR applies to FIA_ATD.1, FIA_UAU.1, and FIA_UID.1.*

Application Note: *The configuration of these parameters is performed with the configuration files /etc/passwd and /etc/shadow which are writeable to the root user only.*

6.1.3.17 Management of TSF data (FMT_MTD.1(SSH))

FMT_MTD.1.1 The TSF shall restrict the ability to **modify the authentication methods provided by the OpenSSH server to the root user.**

Application Note: *This SFR applies to FIA_UAU.5.*

Application Note: *The configuration of this parameter is performed with the configuration file /etc/sshd_config which is writeable to the root user only.*

6.1.3.18 Management of TSF data (FMT_MTD.1(SSL))

FMT_MTD.1.1 The TSF shall restrict the ability to **modify the time interval of user inactivity for locking an interactive session to**

- a) the root user for system wide settings,**
- b) each user for his own sessions, if allowed by the root user.**

Application Note: *This SFR applies to FTA.SSL.1.*

Application Note: *The time interval is configured in /etc/screenrc which is writeable to root only. Normal users can configure the time interval in ~/.screenrc. The screen application enforcing the session locking can be configured to execute with /etc/profile or /etc/login.csh. The root user can place screen execution commands in these Shell startup files that prevent the loading of ~/.screenrc.*

6.1.3.19 Management of TSF data (FMT_MTD.1(VIRT-COMP))

FMT_MTD.1.1 The TSF shall restrict the ability to **modify, delete, initialize the compartment security attributes to the authorized administrator of the virtual machine configuration.**

Application Note: *This SFR applies to FIA_UID.2(VIRT).*

Application Note: *The authorized administrator is the user allowed to access the libvirt daemon and configure virtual machine parameters.*

6.1.3.20 Revocation (FMT_REV.1(OBJ))

FMT_REV.1.1 The TSF shall restrict the ability to revoke **object security attributes defined by SFPs associated with the corresponding object** under the control of the TSF to

- a) DAC permissions: owners of the object and authorized administrator;**
- b) Other security attributes: authorized administrator.**

Application Note: *The privileges that constitute an authorized administrator are defined in the above mentioned FMT_* SFRs which specify the privileges needed to modify object security attributes. The same privileges are required to revoke these security attributes.*

FMT_REV.1.2 The TSF shall enforce the following rules :

- a) The access rights associated with an object shall be enforced when an access check is made;**

b) no specification of other revocation rules.

Application Note: *Revocation of security attributes for named objects imply the revocation of access granted to users other than the owner of the object. Note that the DAC ownership management (which can be also considered as a form of access revocation) is specified in FMT_MSA.1(PSO).*

6.1.3.21 Revocation (FMT_REV.1(USR))

FMT_REV.1.1 The TSF shall restrict the ability to revoke **user security attributes defined by the SFP** associated with the **corresponding user** under the control of the TSF to **authorized administrators** .

Application Note: *The privileges that constitute an authorized administrator are defined in the above mentioned FMT_* SFRs which specify the privileges needed to modify object security attributes. The same privileges are required to revoke these security attributes.*

FMT_REV.1.2 The TSF shall enforce the *following* rules :

- a) The enforcement of the revocation of security-relevant authorizations with the next user-subject binding process during the next authentication of the user;**
- b) No other rules**

Application Note: *The changes are enforced for a new session when the user affected by the change initiates that new session.*

6.1.3.22 Specification of management functions (FMT_SMF.1)

FMT_SMF.1.1 The TSF shall be capable of performing the following management functions:

- a) Management of auditing;**
- b) Management of cryptographic network protocols;**
- c) Management of Persistent Storage Object Access Control Policy;**
- d) Management of Network Information Flow Control Policy;**
- e) Management of identification and authentication policy;**
- f) Management of user security attributes;**
- g) Management of Compartment Access Control Policy;**
- h) Management of Compartment Information Flow Control Policy;**
- i) Management of virtual machine configurations.**

Application Note: *The given list is kept generic intentionally. This ST specifies one iteration of FMT_MTD.1 per management function required by an SFR. For each FMT_MTD.1 iteration, a corresponding application note refers to the covered SFR(s).*

6.1.3.23 Security roles (FMT_SMR.1)

FMT_SMR.1.1 The TSF shall maintain the roles:

- a) User role with the following rights:**
 - i. Users are authorized to modify their own user password;**
 - ii. Users are authorized to modify the access control permissions for the named objects they own;**

- iii. no other rights;
- b) Configurations stored by user space: administrative users defined by the access permissions to the configurations mentioned in the other management SFRs;
- c) Functions provided by the kernel: administrative users defined by capabilities mentioned in other management SFRs.

FMT_SMR.1.2 The TSF shall be able to associate users with roles.

Application Note: Administrative actions can only be performed when the calling subject possesses the above mentioned capabilities which in the TOE configuration is only provided to processes executing with the effective UID or file system UID of zero (also called the root user). As the account for the root user is disabled for direct logon, authorized administrators are defined as users who are assigned to the "wheel" group. This group allows the use of the "su" application which is the only way to assume the root user capabilities.

6.2 Security Functional Requirements Rationale

6.2.1 Coverage

The following table provides a mapping of SFR to the security objectives, showing that each security functional requirement addresses at least one security objective.

Security Functional Requirements	Objectives
FAU_GEN.1	O.AUDITING
FAU_GEN.2	O.AUDITING
FAU_SAR.1	O.AUDITING
FAU_SAR.2	O.AUDITING
FAU_SEL.1	O.AUDITING
FAU_STG.1	O.AUDITING
FAU_STG.3	O.AUDITING
FAU_STG.4	O.AUDITING
FCS_CKM.1(SYM)	O.CRYPTO.NET
FCS_CKM.1(RSA)	O.CRYPTO.NET
FCS_CKM.1(DSA)	O.CRYPTO.NET
FCS_CKM.2(NET)	O.CRYPTO.NET
FCS_CKM.4	O.CRYPTO.NET
FCS_COP.1(NET)	O.CRYPTO.NET
FCS_RNG.1(SSH-DFLT)	O.CRYPTO.NET
FCS_RNG.1(SSH-FIPS)	O.CRYPTO.NET

Security Functional Requirements	Objectives
FDP_ACC.1(PSO)	O.DISCRETIONARY.ACCESS
FDP_ACF.1(PSO)	O.DISCRETIONARY.ACCESS
FDP_IFC.2(NI)	O.NETWORK.FLOW
FDP_IFF.1(NI-IPTables)	O.NETWORK.FLOW
FDP_IFF.1(NI-ebtables)	O.NETWORK.FLOW
FDP_ITC.2(BA)	O.NETWORK.FLOW
FDP_RIP.2	O.AUDITING, O.CRYPTO.NET, O.DISCRETIONARY.ACCESS, O.I&A, O.NETWORK.FLOW
FDP_RIP.3	O.AUDITING, O.CRYPTO.NET, O.DISCRETIONARY.ACCESS, O.I&A, O.NETWORK.FLOW
FIA_AFL.1	O.I&A
FIA_ATD.1(HU)	O.I&A
FIA_ATD.1(TU)	O.NETWORK.FLOW
FIA_SOS.1	O.I&A
FIA_UAU.1	O.I&A
FIA_UAU.5	O.I&A
FIA_UAU.7	O.I&A
FIA_UID.1	O.I&A, O.NETWORK.FLOW
FIA_USB.2	O.I&A
FPT_STM.1	O.AUDITING
FPT_TDC.1(BA)	O.NETWORK.FLOW
FTA_SSL.1	O.I&A
FTA_SSL.2	O.I&A
FTP_ITC.1	O.TRUSTED_CHANNEL
FDP_ACC.2(VIRT)	O.COMP.RESOURCE_ACCESS
FDP_ACF.1(VIRT)	O.COMP.RESOURCE_ACCESS

Security Functional Requirements	Objectives
FDP_ETC.2(VIRT)	O.COMP.INFO_FLOW_CTRL, O.COMP.RESOURCE_ACCESS
FDP_IFC.2(VIRT)	O.COMP.INFO_FLOW_CTRL
FDP_IFF.1(VIRT)	O.COMP.INFO_FLOW_CTRL
FDP_ITC.2(VIRT)	O.COMP.INFO_FLOW_CTRL, O.COMP.RESOURCE_ACCESS
FIA_UID.2(VIRT)	O.COMP.IDENT
FPT_TDC.1(VIRT)	O.COMP.INFO_FLOW_CTRL, O.COMP.RESOURCE_ACCESS
FMT_MSA.1(PSO)	O.MANAGE
FMT_MSA.1(VIRT-CACP)	O.COMP.RESOURCE_ACCESS
FMT_MSA.1(VIRT-CIFCP)	O.COMP.INFO_FLOW_CTRL
FMT_MSA.3(PSO)	O.MANAGE
FMT_MSA.3(NI)	O.MANAGE
FMT_MSA.3(VIRT-CACP)	O.COMP.RESOURCE_ACCESS
FMT_MSA.3(VIRT-CIFCP)	O.COMP.INFO_FLOW_CTRL
FMT_MSA.4(PSO)	O.MANAGE
FMT_MTD.1(AE)	O.MANAGE
FMT_MTD.1(AS)	O.MANAGE
FMT_MTD.1(AT)	O.MANAGE
FMT_MTD.1(AF)	O.MANAGE
FMT_MTD.1(NI)	O.MANAGE
FMT_MTD.1(IAT)	O.MANAGE
FMT_MTD.1(IAF)	O.MANAGE
FMT_MTD.1(IAU)	O.MANAGE
FMT_MTD.1(SSH)	O.MANAGE
FMT_MTD.1(SSL)	O.MANAGE
FMT_MTD.1(VIRT-COMP)	O.COMP.INFO_FLOW_CTRL, O.COMP.RESOURCE_ACCESS
FMT_REV.1(OBJ)	O.MANAGE
FMT_REV.1(USR)	O.MANAGE

Security Functional Requirements	Objectives
FMT_SMF.1	O.MANAGE
FMT_SMR.1	O.MANAGE

Table 8: Mapping of security functional requirements to security objectives

6.2.2 Sufficiency

The following rationale provides justification for each security objective for the TOE, showing that the security functional requirements are suitable to meet and achieve the security objectives:

Security objectives	Rationale
O.AUDITING	<p>The events to be audited are defined in [FAU_GEN.1] and are associated with the identity of the user that caused the event [FAU_GEN.2]. Authorized users are provided the capability to read the audit records [FAU_SAR.1], while all other users are denied access to the audit records [FAU_SAR.2]. The authorized user must have the capability to specify which audit records are generated [FAU_SEL.1]. The TOE prevents the audit log from being modified or deleted [FAU_STG.1] and ensures that the audit log is not lost due to resource shortage [FAU_STG.3, FAU_STG.4]. To support auditing, the TOE is able to maintain proper time stamps [FPT_STM.1].</p> <p>The protection of reused resources ensures that no data leaks from other protected sources [FDP_RIP.2, FDP_RIP.3].</p>
O.CRYPTO.NET	<p>The cryptographically-protected network protocol [FCS_COP.1(NET)] is supported by the generation of symmetric keys [FCS_CKM.1(SYM)], as well as asymmetric keys [FCS_CKM.1(RSA), FCS_CKM.1(DSA)] where the functionality is based on the random number generator as defined by [FCS_RNG.1(SSH-DFLT), FCS_RNG.1(SSH-FIPS)]. As part of the cryptographic network protocol, the TOE securely exchanges the symmetric key with a remote trusted IT system [FCS_CKM.2(NET)]. The TOE ensures that all keys are zeroized upon de-allocation [FCS_CKM.4].</p> <p>The protection of reused resources ensures that no data leaks from other protected sources [FDP_RIP.2, FDP_RIP.3].</p>
O.DISCRETIONARY.ACCESS	<p>The TSF must control access to resources based on the identity of users that are allowed to specify which resources they want to access for storing their data.</p> <p>The access control policy must have a defined scope of control [FDP_ACC.1(PSO)]. The rules for the access control policy are defined [FDP_ACF.1(PSO)].</p> <p>The protection of reused resources ensures that no data leaks from other protected sources [FDP_RIP.2, FDP_RIP.3].</p>
O.NETWORK.FLOW	<p>The network information flow control mechanism controls the information flowing between different entities [FDP_IFC.2(NI)]. The TOE implements a rule-set governing the information flow [FDP_IFF.1(NI-IPTables),</p>

Security objectives	Rationale
	<p>FDP_IFF.1(NI-ebtables)]. To facilitate the information flow control, the information must be identified [FIA_UID.1] based on security attributes the TOE can maintain [FIA_ATD.1(TU)]. The TOE must ensure that security attributes of the network data required by the information flow control policy are correctly interpreted [FDP_ITC.2(BA), FPT_TDC.1(BA)].</p> <p>The protection of reused resources ensures that no data leaks from other protected sources [FDP_RIP.2, FDP_RIP.3].</p>
O.I&A	<p>The TSF must ensure that only authorized users gain access to the TOE and its resources. Users authorized to access the TOE must use an identification and authentication process [FIA_UID.1, FIA_UAU.1]. Multiple I&A mechanisms are allowed as specified in [FIA_UAU.5]. To ensure authorized access to the TOE, authentication data is protected [FIA_ATD.1(HU), FIA_UAU.7]. Proper authorization for subjects acting on behalf of users is also ensured [FIA_USB.2]. The appropriate strength of the authentication mechanism is ensured [FIA_SOS.1]. To support the strength of authentication methods, the TOE is capable of identifying and reacting to unsuccessful authentication attempts [FIA_AFL.1]. In addition, user-initiated and TSF-initiated session locking [FTA_SSL.1, FTA_SSL.2] protect the authenticated user's session.</p> <p>The protection of reused resources ensures that no data leaks from other protected sources [FDP_RIP.2, FDP_RIP.3] are present.</p>
O.MANAGE	<p>The TOE provides management interfaces globally defined in [FMT_SMF.1] for:</p> <ul style="list-style-type: none"> • the access control policies [FMT_MSA.1(PSO), FMT_MSA.3(PSO)]; • the information flow control policy [FMT_MSA.3(NI), FMT_MTD.1(NI)]; • the auditing aspects [FMT_MTD.1(AE), FMT_MTD.1(AS), FMT_MTD.1(AT), FMT_MTD.1(AF)]; • the identification and authentication aspects [FMT_MTD.1(IAT), FMT_MTD.1(IAF), FMT_MTD.1(IAU), FMT_MTD.1(SSH)], • the session locking threshold [FMT_MTD.1(SSL)]. <p>Persistently stored user data is stored either in hierarchical or relational fashion, which implies an inheritance of security attributes from parent object [FMT_MSA.4(PSO)].</p> <p>The rights management for the different management aspects is defined with [FMT_SMR.1].</p> <p>The management interfaces for the revocation of user and object attributes is provided with [FMT_REV.1(OBJ) and FMT_REV.1(USR)].</p>
O.TRUSTED_CHANNEL	<p>The TOE provides a trusted channel protecting communication between a remote trusted IT system and itself [FTP_ITC.1].</p>
O.COMP.INFO_FLOW_CTRL	<p>The information flow control policy covering the runtime of the compartments is specified with [FDP_IFC.2(VIRT)], and [FDP_IFF.1(VIRT)].</p>

Security objectives	Rationale
	<p>As the TOE shall allow export of data belonging to compartments, the TOE assigns the security attributes for enforcing the information flow control policy to the communicated data as specified with [FDP_ETC.2(VIRT)], [FDP_ITC.2(VIRT)], and [FPT_TDC.1(VIRT)].</p> <p>Management of the security attributes for the information flow control policy is specified with [FMT_MSA.1(VIRT-CIFCP)], and [FMT_MSA.3(VIRT-CIFCP)] as well as FMT_MTD.1(VIRT-COMP).</p>
O.COMP.RESOURCE_ACCESS	<p>The access control policy for the resources belonging to the different compartments is defined with [FDP_ACC.2(VIRT)], [FDP_ACF.1(VIRT)].</p> <p>As the TOE shall allow export of data belonging to compartments, the TOE assigns the security attributes for enforcing the access control policy to the communicated data as specified with [FDP_ETC.2(VIRT)], [FDP_ITC.2(VIRT)], and [FPT_TDC.1(VIRT)].</p> <p>Management of the security attributes for the access control policy is specified with [FMT_MSA.1(VIRT-CACP)], and [FMT_MSA.3(VIRT-CACP)] as well as FMT_MTD.1(VIRT-COMP).</p>
O.COMP.IDENT	<p>The identification of compartments to support the information flow control and access control policies is established with [FIA_UID.2(VIRT)].</p>

Table 9: Security objectives for the TOE rationale

6.2.3 Security requirements dependency analysis

The following table demonstrates the dependencies of SFRs modeled in CC Part 2 and how the SFRs for the TOE resolve those dependencies:

Security Functional Requirement	Dependencies	Resolution
FAU_GEN.1	FPT_STM.1	FPT_STM.1
FAU_GEN.2	FAU_GEN.1	FAU_GEN.1
	FIA_UID.1	FIA_UID.1
FAU_SAR.1	FAU_GEN.1	FAU_GEN.1
FAU_SAR.2	FAU_SAR.1	FAU_SAR.1
FAU_SEL.1	FAU_GEN.1	FAU_GEN.1
	FMT_MTD.1	FMT_MTD.1(AE)
FAU_STG.1	FAU_GEN.1	FAU_GEN.1
FAU_STG.3	FAU_STG.1	FAU_STG.1
FAU_STG.4	FAU_STG.1	FAU_STG.1

Security Functional Requirement	Dependencies	Resolution
FCS_CKM.1(SYM)	[FCS_CKM.2 or FCS_COP.1]	FCS_COP.1(NET)
	FCS_CKM.4	FCS_CKM.4
FCS_CKM.1(RSA)	[FCS_CKM.2 or FCS_COP.1]	FCS_COP.1(NET)
	FCS_CKM.4	FCS_CKM.4
FCS_CKM.1(DSA)	[FCS_CKM.2 or FCS_COP.1]	FCS_COP.1(NET)
	FCS_CKM.4	FCS_CKM.4
FCS_CKM.2(NET)	[FDP_ITC.1 or FDP_ITC.2 or FCS_CKM.1]	FCS_CKM.1(SYM) FCS_CKM.1(RSA) FCS_CKM.1(DSA)
	FCS_CKM.4	FCS_CKM.4
FCS_CKM.4	[FDP_ITC.1 or FDP_ITC.2 or FCS_CKM.1]	FCS_CKM.1(SYM)
FCS_COP.1(NET)	[FDP_ITC.1 or FDP_ITC.2 or FCS_CKM.1]	FCS_CKM.1(SYM) FCS_CKM.1(RSA) FCS_CKM.1(DSA)
	FCS_CKM.4	FCS_CKM.4
FCS_RNG.1(SSH-DFLT)	No dependencies.	
FCS_RNG.1(SSH-FIPS)	No dependencies.	
FDP_ACC.1(PSO)	FDP_ACF.1	FDP_ACF.1(PSO)
FDP_ACF.1(PSO)	FDP_ACC.1	FDP_ACC.1(PSO)
	FMT_MSA.3	FMT_MSA.3(PSO)
FDP_IFC.2(NI)	FDP_IFF.1	FDP_IFF.1(NI-IPTables) FDP_IFF.1(NI-ebtables)
FDP_IFF.1(NI-IPTables)	FDP_IFC.1	FDP_IFC.2(NI)
	FMT_MSA.3	FMT_MSA.3(NI)
FDP_IFF.1(NI-ebtables)	FDP_IFC.1	FDP_IFC.2(NI)
	FMT_MSA.3	FMT_MSA.3(NI)
FDP_ITC.2(BA)	[FDP_ACC.1 or FDP_IFC.1]	FDP_IFC.2(NI)
	[FTP_ITC.1 or FTP_TRP.1]	FTP_ITC.1
	FPT_TDC.1	FPT_TDC.1(BA)

Security Functional Requirement	Dependencies	Resolution
FDP_RIP.2	No dependencies.	
FDP_RIP.3	No dependencies.	
FIA_AFL.1	FIA_UAU.1	FIA_UAU.1
FIA_ATD.1(HU)	No dependencies.	
FIA_ATD.1(TU)	No dependencies.	
FIA_SOS.1	No dependencies.	
FIA_UAU.1	FIA_UID.1	FIA_UID.1
FIA_UAU.5	No dependencies.	
FIA_UAU.7	FIA_UAU.1	FIA_UAU.1
FIA_UID.1	No dependencies.	
FIA_USB.2	FIA_ATD.1	FIA_ATD.1(HU)
FPT_STM.1	No dependencies.	
FPT_TDC.1(BA)	No dependencies.	
FTA_SSL.1	FIA_UAU.1	FIA_UAU.1
FTA_SSL.2	FIA_UAU.1	FIA_UAU.1
FTP_ITC.1	No dependencies.	
FDP_ACC.2(VIRT)	FDP_ACF.1	FDP_ACF.1(VIRT)
FDP_ACF.1(VIRT)	FDP_ACC.1	FDP_ACC.2(VIRT)
	FMT_MSA.3	FMT_MSA.3(VIRT-CACP)
FDP_ETC.2(VIRT)	[FDP_ACC.1 or FDP_IFC.1]	FDP_ACC.2(VIRT) FDP_IFC.2(VIRT)
FDP_IFC.2(VIRT)	FDP_IFF.1	FDP_IFF.1(VIRT)
FDP_IFF.1(VIRT)	FDP_IFC.1	FDP_IFC.2(VIRT)
	FMT_MSA.3	FMT_MSA.3(VIRT-CIFCP)
FDP_ITC.2(VIRT)	[FDP_ACC.1 or FDP_IFC.1]	FDP_ACC.2(VIRT) FDP_IFC.2(VIRT)
	[FTP_ITC.1 or FTP_TRP.1]	FTP_ITC.1
	FPT_TDC.1	FPT_TDC.1(VIRT)
FIA_UID.2(VIRT)	No dependencies.	

Security Functional Requirement	Dependencies	Resolution
FPT_TDC.1(VIRT)	No dependencies.	
FMT_MSA.1(PSO)	[FDP_ACC.1 or FDP_IFC.1]	FDP_ACC.1(PSO)
	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MSA.1(VIRT-CACP)	[FDP_ACC.1 or FDP_IFC.1]	FDP_ACC.2(VIRT)
	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MSA.1(VIRT-CIFCP)	[FDP_ACC.1 or FDP_IFC.1]	FDP_IFC.2(VIRT)
	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MSA.3(PSO)	FMT_MSA.1	FMT_MSA.1(PSO)
	FMT_SMR.1	FMT_SMR.1
FMT_MSA.3(NI)	FMT_MSA.1	FMT_MTD.1(NI) is specified to require the management of security attributes for the Network Information Flow Control Policy, just as a potential FMT_MSA.1(NI) would have been specified. However, the Network Information Flow Control Policy is not required to be enforced when managing the security attributes, as the management aspect of the network information flow control functionality is not protected by the network information flow control mechanism. Therefore, FMT_MSA.1 is not applicable and is replaced with FMT_MTD.1(NI).
	FMT_SMR.1	FMT_SMR.1
FMT_MSA.3(VIRT-CACP)	FMT_MSA.1	FMT_MSA.1(VIRT-CACP)
	FMT_SMR.1	FMT_SMR.1
FMT_MSA.3(VIRT-CIFCP)	FMT_MSA.1	FMT_MSA.1(VIRT-CIFCP)
	FMT_SMR.1	FMT_SMR.1
FMT_MSA.4(PSO)	[FDP_ACC.1 or FDP_IFC.1]	FDP_ACC.1(PSO)
FMT_MTD.1(AE)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1

Security Functional Requirement	Dependencies	Resolution
FMT_MTD.1(AS)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MTD.1(AT)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MTD.1(AF)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MTD.1(NI)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MTD.1(IAT)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MTD.1(IAF)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MTD.1(IAU)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MTD.1(SSH)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MTD.1(SSL)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_MTD.1(VIRT-COMP)	FMT_SMR.1	FMT_SMR.1
	FMT_SMF.1	FMT_SMF.1
FMT_REV.1(OBJ)	FMT_SMR.1	FMT_SMR.1
FMT_REV.1(USR)	FMT_SMR.1	FMT_SMR.1
FMT_SMF.1	No dependencies.	
FMT_SMR.1	FIA_UID.1	FIA_UID.1

Table 10: TOE SFR dependency analysis

6.3 Security Assurance Requirements

The security assurance requirements for the TOE are the Evaluation Assurance Level 4 components as specified in [CC] part 3, augmented by ALC_FLR.3.

The following table shows the Security assurance requirements, and the operations performed on the components according to CC part 3: iteration (Iter.), refinement (Ref.), assignment (Ass.) and selection (Sel.).

Security assurance class	Security assurance requirement	Source	Operations			
			Iter.	Ref.	Ass.	Sel.
ADV Development	ADV_ARC.1 Security architecture description	CC Part 3	No	No	No	No
	ADV_FSP.4 Complete functional specification	CC Part 3	No	No	No	No
	ADV_IMP.1 Implementation representation of the TSF	CC Part 3	No	No	No	No
	ADV_TDS.3 Basic modular design	CC Part 3	No	No	No	No
AGD Guidance documents	AGD_OPE.1 Operational user guidance	CC Part 3	No	No	No	No
	AGD_PRE.1 Preparative procedures	CC Part 3	No	No	No	No
ALC Life-cycle support	ALC_CMC.4 Production support, acceptance procedures and automation	CC Part 3	No	No	No	No
	ALC_CMS.4 Problem tracking CM coverage	CC Part 3	No	No	No	No
	ALC_DEL.1 Delivery procedures	CC Part 3	No	No	No	No
	ALC_DVS.1 Identification of security measures	CC Part 3	No	No	No	No
	ALC_FLR.3 Systematic flaw remediation	CC Part 3	No	No	No	No
	ALC_LCD.1 Developer defined life-cycle model	CC Part 3	No	No	No	No
	ALC_TAT.1 Well-defined development tools	CC Part 3	No	No	No	No
ASE Security Target evaluation	ASE_INT.1 ST introduction	CC Part 3	No	No	No	No
	ASE_CCL.1 Conformance claims	CC Part 3	No	No	No	No
	ASE_SPD.1 Security problem definition	CC Part 3	No	No	No	No
	ASE_OBJ.2 Security objectives	CC Part 3	No	No	No	No
	ASE_ECD.1 Extended components definition	CC Part 3	No	No	No	No
	ASE_REQ.2 Derived security requirements	CC Part 3	No	No	No	No
	ASE_TSS.1 TOE summary specification	CC Part 3	No	No	No	No
ATE Tests	ATE_COV.2 Analysis of coverage	CC Part 3	No	No	No	No
	ATE_DPT.1 Testing: basic design	CC Part 3	No	No	No	No
	ATE_FUN.1 Functional testing	CC Part 3	No	No	No	No
	ATE_IND.2 Independent testing - sample	CC Part 3	No	No	No	No

Security assurance class	Security assurance requirement	Source	Operations			
			Iter.	Ref.	Ass.	Sel.
AVA Vulnerability assessment	AVA_VAN.3 Focused vulnerability analysis	CC Part 3	No	No	No	No

Table 11: Security assurance requirements

6.4 Security Assurance Requirements Rationale

The basis for the justification of EAL4 augmented with ALC_FLR.3 is the threat environment experienced by the typical consumers of the TOE. This matches the package description for EAL4 (enhanced-basic).

7 TOE Summary Specification

7.1 TOE Security Functionality

The following section explains how the security functions are implemented. The different TOE security functions cover the various SFR classes.

The primary security features of the TOE are:

- Audit
- Cryptographically secured communication channels
- Packet filter
- Identification and Authentication
- Discretionary Access Control
- Authoritative Access Control
- Virtual machine environments
- Security Management

7.1.1 Audit

The Lightweight Audit Framework (LAF) is designed to be an audit system for Linux compliant with the requirements from Common Criteria. LAF is able to intercept all system calls as well as retrieving audit log entries from privileged user space applications. The subsystem allows configuring the events to be actually audited from the set of all events that are possible to be audited. Those events are configured in a specific configuration file and then the kernel is notified to build its own internal structure for the events to be audited.

7.1.1.1 Audit functionality

The Linux kernel implements the core of the LAF functionality. It gathers all audit events, analyzes these events based on the audit rules and forwards the audit events that are requested to be audited to the audit daemon executing in user space.

Audit events are generated in various places of the kernel. In addition, a user space application can create audit records which needs to be fed to the kernel for further processing.

The audit functionality of the Linux kernel is configured by user space applications which communicate with the kernel using a specific netlink communication channel. This netlink channel is also to be used by applications that want to send an audit event to the kernel.

The kernel netlink interface is usable only by applications possessing the following capabilities:

- CAP_AUDIT_CONTROL: Performing management operations like adding or deleting audit rules, setting or getting auditing parameters;
- CAP_AUDIT_WRITE: Submitting audit records to the kernel which in turn forwards the audit records to the audit daemon.

Based on the audit rules, the kernel decides whether an audit event is discarded or to be sent to the user space audit daemon for storing it in the audit trail. The kernel sends the message to the audit daemon again using the above mentioned netlink communication channel. The audit daemon writes the audit records to the audit trail. An internal queuing mechanism is used for this purpose. When the queue does not have sufficient space to hold an audit record the TOE switches into single user mode, is halted or the audit daemon executes an administrator-specified notification action

depending on the configuration of the audit daemon. This ensures that audit records do not get lost due to resource shortage and the administrator can backup and clear the audit trail to free disk space for new audit logs.

Access to audit data by normal users is prohibited by the discretionary access control function of the TOE, which is used to restrict the access to the audit trail and audit configuration files to the system administrator only.

The system administrator can define the events to be audited from the overall events that the Lightweight Audit Framework using simple filter expressions. This allows for a flexible definition of the events to be audited and the conditions under which events are audited. The system administrator is also able to define a set of user IDs for which auditing is active or alternatively a set of user IDs that are not audited.

The system administrator can select files to be audited by adding them to a watch list that is loaded into the kernel.

This security function covers the SFRs of: FAU_GEN.1, FAU_GEN.2, FAU_SAR.2, FAU_SEL.1, FAU_STG.1, FAU_STG.3, FAU_STG.4, FPT_STM.1.

7.1.1.2 Audit trail

An audit record consists of one or more lines of text containing fields in a “keyword=value” tagged format. The following information is contained in all audit record lines:

- Type: indicates the source of the event, such as SYSCALL, FS_WATCH, USER, or LOGIN
- Timestamp: Date and time the audit record was generated
- Audit ID: unique numerical event identifier
- Login ID (“audit”), the user ID of the user authenticated by the system (regardless if the user has changed his real and / or effective user ID afterwards)
- Effective user ID: the effective user ID of the process at the time the audit event was generated
- Success or failure (where appropriate)
- SELinux label of the subject that caused the event
- Process ID of the subject that caused the event (PID)

This information is followed by event specific data. In some cases, such as SYSCALL event records involving file system objects, multiple text lines will be generated for a single event, these all have the same time stamp and audit ID to permit easy correlation.

The audit trail is stored in ASCII text. The TOE provides tools for managing ASCII files that can be used for post-processing of audit data. These tools include:

- less - reads the ASCII audit data
- ausearch - allows selective extraction of records from the audit trail using defined selection criteria
- sort - The audit records are listed in chronological order by default. The sort utility can be used together with ausearch to use a different sorting order.

The audit trail is stored in files which are accessible by root only.

This security function covers the SFRs of: FAU_SAR.1.

7.1.2 Cryptographic services

The TOE provides cryptographically secured network communication channels to allow remote users to interact with the TOE. Using one of the following cryptographically secured network channels, a user can request the following services:

- **OpenSSH:** The OpenSSH application provides access to the command line interface of the TOE. Users may employ OpenSSH for interactive sessions as well as for non-interactive sessions. The console provided via OpenSSH provides the same environment as a local console. OpenSSH implements the SSHv2 protocol.
- **libvirt:** The libvirt daemon is the management facility to allow remote users to configure virtual machines. The configuration covers all aspects such as assigning of resources, starting or stopping of virtual machines. libvirt directly interacts with the virtual machines. This interface can be protected using OpenSSH.

The cryptographic primitives for implementing the above mentioned cryptographic communication protocols are provided by OpenSSL.

7.1.2.1 SSHv2 Protocol

The TOE provides the Secure Shell Protocol Version 2 (SSH v2.0) to allow users from a remote host to establish a secure connection and perform a logon to the TOE.

The following table documents implementation details concerning the OpenSSH implementation's compliance to the relevant standards. It addresses areas where the standards permit different implementation choices such as optional features.

Reference	Description	Implementation Details
RFC 4253 chapter 5	Compatibility with old SSH versions	The OpenSSH implementation is capable of interoperating with clients and servers using the old 1.x protocol. That functionality is explicitly disabled in the evaluated configuration, it permits protocol version 2.0 exclusively.
RFC 4253 section 6.2	Compression	OpenSSH supports the OPTIONAL "zlib" compression method.
RFC 4253 section 6.3	Encryption	The ciphers supported in the evaluated configuration are listed in FCS_COP.1(NET) for the SSH protocol.
RFC 4252 chapter 7	Public Key Authentication Method: "publickey"	This REQUIRED authentication method is supported by OpenSSH but can be disabled by the administrator of the OpenSSH daemon.
RFC 4252 chapter 8	Password Authentication Method: "password"	This SHOULD authentication method is supported by OpenSSH but can be disabled by the administrator of the OpenSSH daemon.
RFC 4252 chapter 8	Password change request and setting new password	The OpenSSH implementation supports the optional password change mechanism in the evaluated configuration.
RFC 4252 chapter 9	Host-Based Authentication: "hostbased"	This OPTIONAL authentication method is disabled in the evaluated configuration.

Table 12: SSH implementation notes

The TOE supports the generation of RSA and DSA key pairs. These key pairs are used by OpenSSH for the host keys as well as for the per-user keys. When a user registers his public key with the user he wants to access on the server side, a key-based authentication can be performed instead of a password-based authentication. The key generation mechanism uses the Linux kernel random number generator. The evaluated configuration permits the import of externally-generated RSA/DSA key pairs.

This security function covers the SFRs of: FCS_CKM.1(RSA), FCS_CKM.1(DSA), FCS_RNG.1(SSH-DFLT), FCS_RNG.1(SSH-FIPS).

The TOE supports the following security functions of the SSH v2.0 protocol:

- Establishing a secure communication channel using the following cryptographic functions provided by the SSH v2.0 protocol:
 - Encryption as defined in section 4.3 of RFC 4253 - the keys are generated using the Linux kernel random number generator;
 - Diffie-Hellman key exchange as defined in section 6.1 of RFC 4253;
 - The keyed hash function for integrity protection as defined in section 4.4 of RFC 4253.

Note: The protocol supports more cryptographic algorithms than the ones listed above. Those other algorithms are not covered by this evaluation and should be disabled or not used when running the evaluated configuration.

- Performing user authentication using the standard password-based authentication method the TOE provides for users (password authentication method as defined in chapter 5 of RFC 4252).
- Performing user authentication using a RSA or DSA key-based authentication method (public key authentication method as defined in chapter 5 of RFC 4252).
- Checking the integrity of the messages exchanged and close down the connection in case an integrity error is detected.

The OpenSSH applications of `sshd`, `ssh` and `ssh-keygen` use the OpenSSL random number generator seeded by `/dev/random` to generate cryptographic keys. OpenSSL provides different DRNGs depending whether the FIPS 140-2 mode is enabled in the system.

This security function covers the SFRs of: FCS_CKM.1(SYM), FCS_CKM.2(NET), FCS_CKM.4, FCS_COP.1(NET), FCS_RNG.1(SSH-DFLT), FCS_RNG.1(SSH-FIPS), FTP_ITC.1.

7.1.3 Packet filter

The Linux kernel's network stack implementation follows the layering structure of the network protocols. It implements the code for handling the link layer as well as the network layer. For those layers, independent filter mechanism are provided:

- Link layer: `etables` implements the filtering mechanism for bridges
- Network layer: `netfilter/iptables` implements the filtering mechanism for non-bridge interfaces

7.1.3.1 Network layer filtering

Netfilter

Netfilter is a framework for packet mangling, implemented in the Linux kernel network stack handling the network layer. The netfilter framework comprises of the following parts:

- The IP stack defines five hooks which are well-defined points in a network packet's traversal of the IP protocol stack. Each of the hooks, the network stack will call the netfilter framework allowing it to operate on the entire packet. (Note: the netfilter framework provides such hooks in a number of network protocol implementations, but the TOE only supports IP as outlined above. Therefore, the ST specification only covers the IP protocol.)
- The netfilter framework provides register functions for other kernel parts to listen to the different hooks. When a packet traverses one of the hooks and passed to the netfilter framework, it invokes every registered kernel part. These kernel parts then can examine the packet and possible alter it. As part of the examination, these kernel parts can instruct the netfilter framework to discard the packet, to allow it to pass, or to queue it to user space.
- When a packet is marked to be queued to user space, the netfilter framework handles the asynchronous communication with user space.

The netfilter framework implements the five hooks at the following points in the packet traversal chain:

- When the packet enters the network layer of the TOE and after applying some sanity checks, but before the routing table is consulted, the `NF_IP_PRE_ROUTING` hook is triggered.
- After passing the routing table decision and the routing code marks the packet to be targeted for another host, the `NF_IP_FORWARD` hook is triggered.
- After passing the routing table decision and the routing code marks the packet to be targeted for the local system, the `NF_IP_LOCAL_IN` hook is triggered.
- When the packet traversed all of the network stack and is about to be placed on the wire again, the `NF_IP_POST_ROUTING` hook is triggered.
- When a packet is generated locally, the `NF_IP_LOCAL_OUT` hook is triggered before the routing table is consulted.

IPTables

All communication on the network layer can be controlled by the IPTables framework.

The TOE implements a packet filter as part of the network stack provided with the Linux kernel. The combination of IPTables and netfilter implements the packet filter which provides stateful and stateless packet filtering for network communication by inspecting the IP header, the TCP header, UDP header and/or ICMP header of every network packet that passes the network stack.

The packet selection system called IP Tables uses the netfilter framework to implement the actual packet filtering logic on the network layer for the TCP/IP protocol family.

Note: IPTables is able to perform Network Address Translation (NAT) as well as Port Address Translation (PAT) for simple as well as more complex protocols. This mechanism is out of scope for the evaluation. Furthermore, packet mangling support is provided with IPTables which is also out of scope for the evaluation.

IPTables registers all hooks provided by the netfilter framework. The NAT/PAT mechanism uses the pre-routing and post-routing hooks whereas the packet filtering capability is enforced on the local-in, local-out and forwarding hooks.

IPTables consists of the following two components:

- In-kernel packet filter enforcement: The kernel-side of IPTables use the netfilter framework as indicated above. Three lists of packet filter rules are enforced by the kernel mechanism: one for each netfilter framework hook that applies to packet filtering. When a packet is analyzed by the IPTables kernel modules, they first select the applicable list based on the

hook where the netfilter framework triggered IPTables. Each list contains zero or more rules which are iterated sequentially. A rule consists of a matching part (also called the "match extension") and an action part (also called the "target extension"). When a rule is applied to a packet, the kernel modules first applies the matching part of the rule. If the packet matches, the action part is enforced. If the action part contains a decision of the fate of the packet (to accept it, to drop it, or to drop it and sending a notification to the sender), the rule list validation stops for this packet. If the action part contains a modification instruction or log instruction for the packet, the rule list validation continues after performing this operation. When the rule list is iterated through and a packet could not be matched by a rule with a decision action (accept, drop), the default decision action applicable to the list is enforced. This default action is either to accept the packet, to drop the packet, or to drop the packet and send a notification to the sender.

- User space configuration application: The user space application iptables(1) allows the configuration of the IPTables kernel components. The application allows the specification of one rule per invocation where a rule contains the above mentioned matching part and action part. The tool also allows modification or deletion of existing rules as well as configuration of the default action. When using the tool, each invocation must specify the netfilter framework hook to which the rule applies to. See the man page of iptables(1) for more details.

7.1.3.2 Link layer filtering

eatables chains

Similarly to the netfilter hooks in the network layer handling protocol, the link layer code handling the bridging functionality implements chains that are used by eatables to apply filtering.

The netfilter framework is used to implement the eatables chains at the following points in the packet traversal logic:

- When the packet enters the link layer of the TOE and after applying some sanity checks, but before the bridging decision made, the BROUTING chain is triggered.
- After passing the BROUTING chain but still before the bridging decision, the PREROUTING chain is triggered. When the bridging code decides that the frame is not intended for a bridge, it is forwarded to the network layer and processing on the link layer stops.
- After passing the bridge routing decision and the routing code marks the packet to be targeted for another host, the FORWARD chain is triggered.
- After passing the bridge routing decision and the routing code marks the packet to be targeted for the local system, the INPUT chain is triggered.
- When a packet is generated locally and the bridging decision marks the frame to be intended for a bridge, the OUTPUT chain is triggered.
- When the packet traversed all of the bridge logic and is about to be placed on the wire again, the POSTROUTING chain is triggered.

eatables filtering rules

The packet selection system called eatables uses the netfilter framework to implement the actual packet filtering logic for Ethernet frames routed through bridges.

Bridged networking communication is only visible on the link layer. As the host system uses bridges to connect virtual machines to the environment, only ebtables can be used to control the traffic flow. The data encapsulated within the Ethernet frames are not routed through the host system's network stack and can therefore not be analyzed and controlled by the IPTables framework outlined above.

Even though the network stack of the host system does not analyze the bridged Ethernet frames any more, the ebtables framework receives the entire frame and is allowed to operate on that frame. Therefore, ebtables can analyze the protocol inside the Ethernet frame.

Please note that the ST author is perfectly aware that the term "packet" is used for the network layer (which includes the TCP/IP protocol family), whereas the term "frame" applies to the link layer (which includes the Ethernet frames). However, when speaking about ebtables that can analyze IP packets inside the link layer, this ST refers to them as packets/frames.

The ebtables framework can apply filters based on the following concepts:

- Matching of the frame/packet: Matches can be configured individually for different protocols. For IP, the source and destination address, the protocol as well as the source and destination port can be specified. In addition, the interfaces the frame entered or will leave the system can be used for matching.
- Action to be performed on matched frames: Similarly to IPTables, ebtables must be configured to perform an action on the matched frame. The action can either be to discard the packet or to accept it.

This security function covers the SFRs of:

- Packet filtering rules: FDP_IFC.2(NI), FDP_IFF.1(NI-*)
- Interpretation of network protocol: FIA_UID.1, FDP_ITC.2, FPT_TDC.1(BA)
- Maintenance of rules: FIA_ATD.1(TU)

7.1.4 Identification and Authentication

User identification and authentication in the TOE includes all forms of interactive login (e.g. using the SSH protocol or log in at the local console) as well as identity changes through the su command. These all rely on explicit authentication information provided interactively by a user. In addition, the key-based authentication mechanism of the OpenSSH server is another form of authentication.

7.1.4.1 PAM-based identification and authentication mechanisms

Linux uses a suite of libraries called the "Pluggable Authentication Modules" (PAM) that allow an administrative user to choose how PAM-aware applications authenticate users. The TOE provides PAM modules that implement all the security functionality to:

- Provides login control and establishing all UIDs, GIDs and login ID for a subject
- Ensure the quality of passwords
- Enforce limits for accounts (such as the number of maximum concurrent sessions allowed for a user)
- Enforce the change of passwords after a configured time including the password quality enforcement
- Enforcement of locking of accounts after failed login attempts
- Restriction of the use of the root account to certain terminals
- Restriction of the use of the su command

- Set up of the sensitivity label and file system name space

The login processing sets the real, file system effective and login UID as well as the real, effective, file system GID and the set of supplemental GIDs of the subject that is created. It is of course up to the client application usually provided by a remote system to protect the user's entry of a password correctly (e. g. provide only obscured feedback).

During login processing, the user is shown a banner. After successful authentication, the login time is recorded.

After a successful identification and authentication, the TOE initiates a session for the user and spawns the initial login shell as the first process the user can interact with. The TOE provides a mechanism to lock a session either automatically after a configurable period of inactivity for that session or upon the user's request.

This security function covers the SFRs of: FDP_RIP.3, FIA_AFL.1, FIA_SOS.1, FI_UAU.1, FIA_UID.1, FIA_UAU.5, FIA_UAU.7, FIA_USB.2.

7.1.4.2 User Identity Changing

Users can change their identity (i.e., switch to another identity) using the su command. When switching identities, the real, file system and effective user ID and real, file system and effective group ID are changed to the one of the user specified in the command (after successful authentication as this user).

The primary use of the su command within the TOE is to allow appropriately authorized individuals the ability to assume the root identity to perform administrative actions. In this system the capability to login as the root identity has been restricted to defined terminals only. In addition the use of the su command to switch to root has been restricted to users belonging to a special group. Users that don't have access to a terminal where root login is allowed and are not member of that special group will not be able to switch their real, file system and effective user ID to root even if they would know the authentication information for root. Note that when a user executes a program that has the setuid bit set, only the effective user ID and file system ID are changed to that of the owner of the file containing the program while the real user ID remains that of the caller. The login ID is neither changed by the su command nor by executing a program that has the setuid or setgid bit set as it is used for auditing purposes.

Note: The login ID is not retained for the following special case:

1. User A logs into the system.
2. User A uses su to change to user B.
3. User B now edits the cron or at job queue to add new jobs. This operation is appropriately audited with the proper login ID.
4. Now when the new jobs are executed as user B, the system does not provide the audit information that the jobs are created by user A.

The su command invokes the common authentication mechanism to validate the supplied authentication.

This security function covers the SFRs of: FIA_USB.2.

7.1.4.3 Authentication Data Management

Each TOE instance maintains its own set of users with their passwords and attributes. Although the same human user may have accounts on different servers interconnected by a network and running an instantiation of the TOE, those accounts and their parameter are not synchronized on different TOE instances. As a result the same user may have different user names, different user Ids, different passwords and different attributes on different machines within the networked environment. Existing mechanism for synchronizing this within the whole networked system are not subject to this evaluation.

Each TOE instance within the network maintains its own administrative database by making all administrative changes on the local TOE instance. System administration has to ensure that all machines within the network are configured in accordance with the requirements defined in this Security Target.

The file `/etc/passwd` contains for each user the user's name, the id of the user, an indicator whether the password of the user is valid, the principal group id of the user and other (not security relevant) information. The file `/etc/shadow` contains for each user a hash of the user's password, the userid, the time the password was last changed, the expiration time as well as the validity period of the password and some other information that are not subject to the security functions as defined in this Security Target. Users are allowed to change their passwords by using the `passwd` command. This application is able to read and modify the contents of `/etc/shadow` for the user's password entry, which would ordinarily be inaccessible to a non-privileged user process. Users are also warned to change their passwords at login time if the password will expire soon, and are prevented from logging in if the password has expired.

The time of the last successful logins is recorded in the `/var/log/tallylog` file.

The TOE displays informative banners before or during the login to users. The banners can be specified with the files `/etc/issue` for log ins via the physical console or `/etc/issue.net` for remote log ins, such as via SSH. When performing a log in on the physical console, the banner is displayed above the username and password prompt. For log ins via SSH, the banner is displayed to the remote peer during the SSH-session handshake takes place. The remote SSH client will display the banner to the user. When using the provided OpenSSH client, the banner is displayed when the user instructs the OpenSSH client to log into the remote system.

This security function covers the SFRs of: FIA_ATD.1(HU) and FMT_REV.1(USR).

7.1.4.4 SSH key-based authentication

In addition to the PAM-based authentication outlined above, the OpenSSH server is able to perform a key-based authentication. When a user wants to log on, instead of providing a password, the user applies his SSH key. After a successful verification, the OpenSSH server considers the user as authenticated and performs the PAM-based operations as outlined above.

To establish a key-based authentication, a user first has to generate an RSA or DSA key pair. The private part of the key pair remains on the client side. The public part is copied to the server into the file `.ssh/authorized_keys` which resides in the home directory of the user he wants to log on as. When the login operation is performed the SSHv2 protocol tries to perform the "publickey" authentication using the private key on the client side and the public key found on the server side. The operations performed during the publickey authentication is defined in RFC 4252 chapter 7.

Users have to protect their private key part the same way as protecting a password. Appropriate permission settings on the file holding the private key is necessary. To strengthen the protection of the private key, the user can encrypt the key where a password serves as key for the encryption operation. See `ssh-keygen(1)` for more information.

This security function covers the SFRs of: FI_UAU.1, FIA_UID.1, FIA_UAU.5, FIA_SOS.1.

7.1.4.5 Session locking

The TOE uses the `screen(1)` application which locks the current session of the user either after an administrator-specified time of inactivity or upon the user's request.

To unlock the session, the user must supply his password. Screen uses PAM to validate the password and allows the user to access his session after a successful validation.

This security function covers the SFRs of: FTA_SSL.1, FTA_SSL.2.

7.1.5 Discretionary Access Control

The general policy enforced is that subjects (i.e., processes) are allowed only the accesses specified by the policies applicable to the object the subject requests access to. Further, the ability to propagate access permissions is limited to those subjects who have that permission, as determined by the policies applicable to the object the subject requests access to.

A subject may possess one or more of the following capabilities which provide the following exemptions from the DAC mechanism:

- **CAP_DAC_OVERRIDE**: A process with this capability is exempt from all restrictions of the discretionary access control and can perform any action desired. For the execution of a file, the permission bit vector of that file must contain at least one execute bit.
- **CAP_DAC_READ_SEARCH**: A process with this capability overrides all DAC restrictions regarding read and search on files and directories.

DAC provides the mechanism that allows users to specify and control access to objects that they own. DAC attributes are assigned to objects at creation time and remain in effect until the object is destroyed or the object attributes are changed. DAC attributes exist for, and are particular to, each type of named object known to the TOE. DAC is implemented with permission bits and, when specified, ACLs.

The outlined DAC mechanism applies only to named objects which can be used to store or transmit user data. Other named objects are also covered by the DAC mechanism but may be supplemented by further restrictions. These additional restrictions are out of scope for this evaluation. Examples of objects which are accessible to users that cannot be used to store or transmit user data are: virtual file systems externalizing kernel data structures (such as most of `procfs`, `sysfs`, `binfmt_misc`) and process signals.

During creation of objects, the TSF ensures that all residual contents is removed from that object before making it accessible to the subject requesting the creation.

When data is imported into the TOE (such as when mounting disks created by other trusted systems), the TOE enforces the permission bits and ACLs applied to the file system objects.

7.1.5.1 Permission bits

The TOE supports standard UNIX permission bits to provide one form of DAC for file system objects in all supported file systems. There are three sets of three bits that define access for three categories of users: the owning user, users in the owning group, and other users. The three bits in each set indicate the access permissions granted to each user category: one bit for read (r), one for write (w) and one for execute (x). Note that write access to file systems mounted as read only (e. g. CD-ROM) is always rejected (the exceptions are character and block device files which can still be written to as write operations do not modify the information on the storage media). Also, write access to file system objects marked as immutable is always rejected. The SAVETXT attribute is used for world-writable temp directories preventing the removal of files by users other than the owner.

Each process has an inheritable "umask" attribute which is used to determine the default access permissions for new objects. It is a bit mask of the user/group/other read/write/execute bits, and specifies the access bits to be removed from new objects. For example, setting the umask to "002" ensures that new objects will be writable by the owner and group, but not by others. The umask is defined by the administrator in the /etc/login.defs file or 022 by default if not specified.

This security function covers the SFRs of FDP_ACC.1(PSO), FDP_ACF.1(PSO), FDP_ITC.1(BA), FDP_RIP.2, FPT_TDC.1(BA), FMT_REV.1(OBJ).

7.1.5.2 Access Control Lists (ACLs)

The TOE provides support for POSIX type ACLs to define a fine grained access control on a user basis. ACLs are supported for all file system objects stored with the following file systems:

- ext3
- tmpfs

An ACL entry contains the following information:

- A tag type that specifies the type of the ACL entry
- A qualifier that specifies an instance of an ACL entry type
- A permission set that specifies the discretionary access rights for processes identified by the tag type and qualifier

An ACL contains exactly one entry of three different tag types (called the „required ACL entries“ forming the "minimum ACL"). The standard UNIX file permission bits as described in the previous section are represented by the entries in the minimum ACL.

A default ACL is an additional ACL which may be associated with a directory. This default ACL has no effect on the access to this directory. Instead the default ACL is used to initialize the ACL for any file that is created in this directory. If the new file created is a directory it inherits the default ACL from its parent directory. When an object is created within a directory and the ACL is not defined with the function creating the object, the new object inherits the default ACL of its parent directory as its initial ACL.

7.1.5.3 File system objects

Access to file system objects is generally governed by permission bits. For the ext3 file system, ACLs are supported.

File system objects access checks are performed when the object is initially opened, and are not checked on each subsequent access. Changes to access controls (i.e., revocation) are effective with the next attempt to open the object.

7.1.6 Authoritative Access Control

The TOE supports authoritative access control using sensitivity labels automatically attached to processes and objects. This policy is enforced by the SELinux security module and the TOE specific SELinux policy.

The name "authoritative" access control is based on the fact that only authorized administrators can set and modify any labels assigned to objects or subject. Therefore, the access control is not at the discretion of a user like for DAC.

SELinux together together with its policy is used to implement the following policies:

- Multi-category system (MCS): The MCS policy is applied to virtual machines and their resources.

7.1.6.1 Resource access control for virtual machines

The TOE implements the following types of access control restrictions to limit virtual machines to access only their resources:

- SELinux-based: An SELinux policy is provided that implements the multi-category system. The non-hierarchical categories maintained with the SELinux label are used to label virtual machines and their resources. The libvirt virtual machine management daemon assigns a unique category to each virtual machine and its resources. The access control rules only allow subjects (i.e. virtual machines) access to virtual machine resources when the category of the subject is identical to the category of the object.

This security function covers the SFRs of: FDP_ACC.2(VIRT), FDP_ACF.1(VIRT).

7.1.7 Virtual machine environments


KVM is implemented as part of the Linux kernel supported by user space code. It consists of two essential components that implement VMM functionality: the KVM Linux kernel module and QEMU for hardware emulation. The use of QEMU implies that KVM provides full virtualization to its guests and can, therefore, execute unaltered guest operating systems.

The KVM Linux kernel module implements memory management and virtual machine maintenance functionality. This kernel extension makes the entire Linux kernel the hypervisor. Virtual machines are treated by the Linux kernel as normal applications. The kernel schedules them like applications, and they can be handled like applications. As such, the process implementing a virtual machine can be seen in process listings and it can be sent regular signals, like SIGTERM.

The figure showing the virtual machine handling depicts that from the Linux kernel perspective, the virtual machine is just another process. However, the virtual machine process has a special layout. As depicted in the figure, the process image is split into two parts. The first part hosts a regular application logic executing in user mode (the white part of the application box in the figure) – this is used to maintain the QEMU I/O virtualization and some other small KVM-related software. The second part contains the image of the guest code, usually an operating system (the gray part of the application box), which executes in guest mode. This implies that the entire memory used

for the guest operating system is allocated by the QEMU application. The kernel keeps track of which parts of the application belong to the guest operating system and which parts to the regular application.

Figure 1: KVM virtual machine handling by the Linux kernel



When the kernel releases control of the CPU to the virtual machine process, it sets the processor state of the CPU to the user state when calling the regular application logic in user mode. However, when returning control of the CPU to the guest code, the CPU can be set either to supervisor state or user state, depending on the state of the CPU when the Linux kernel initially obtained control.

The figure showing the virtual machine handling also illustrates the overall logic flow that connects the Linux kernel with the regular application logic and the guest operating system. In essence, this logic flow is an endless loop, which is performed as follows:

1. The regular application logic executing in user mode sets up the virtual machine configuration by instructing the kernel to allocate memory for the application, CPU and other resources. The kernel sets up these resources and assigns them to the calling process. After setup is complete, the kernel is instructed to execute the guest. This phase starts the loop and is not executed again during the loop.
2. The kernel now causes the hardware to enter guest mode. If the processor exits guest mode due to an event such as an external interrupt or a shadow page table fault, the kernel performs the necessary handling and resumes guest execution. If the exit is due to an I/O instruction or a signal queued to the process, then the kernel exits to the regular application logic in user mode.
3. The processor executes guest code until it encounters an instruction that needs assistance, a fault, or an external interrupt. The processor then returns control to the VMM kernel.
4. If the kernel detects an exit of the guest code due to an I/O instruction or a signal, or until an external event such as arrival of a network packet or a timeout occurs, the kernel invokes the user mode component of the virtual machine process. The processed I/O instructions cover programmed I/O (PIO) whose implementation is not as complex as the second set of processed I/O instructions, the memory mapped I/O (MMIO). QEMU and a small extension for making QEMU KVM-aware is used to implement the I/O handling as it implements a number of emulated devices and mediates access to real resources when a device is accessed via the I/O instruction. QEMU may alter the virtual processor state to reflect the emulated I/O instruction result to the calling guest code. The modification of the virtual processor state is one with IOCTLs to a file descriptor that QEMU has allocated when setting up the virtual machine. This file descriptor is used to store all virtual machine and virtual CPU data relevant for executing the virtual machine. As the file descriptor is bound to one process only, the kernel implicitly ensures that a QEMU instance can only operate on its virtual machine and virtual CPUs. Once QEMU completes the I/O operation, it signals the kernel that the guest code can resume execution which is implemented with step 2 above.

In this architecture, regular applications (i.e., applications executing only in user state of the CPU and having full access to all services of the Linux kernel and, therefore, other parts of the operating system) coexist with applications that host virtual machines.

The libvirt management daemon sets up virtual machines and controls the resources assigned to virtual machines. To support the separation of virtual machines, libvirt uses the following capabilities:

- Every virtual machine process executes with the normal, unprivileged user ID of “qemu” and the group ID of “qemu”. This implies that these processes do not possess root privilege.
- Libvirt sets up the unique category for a virtual machine and assigns it to the virtual machine and its resources. The resource access control functionality defined as an independent security functionality.

Virtual machines are associated with one or more unique IP addresses that can be used to communicate with other virtual machines on the same host or with other external entities. The TOE ensures that the configured IP addresses are used by the virtual machines for any network-related communication.

This security function for virtual machine maintenance covers the SFRs FDP_IFC.2(VIRT), FDP_IFF.1(VIRT), FIA_UID.2(VIRT).

This security function for virtual machine communication covers all SFRs of FDP_ETC.2(VIRT), FDP_ITC.2(VIRT), FDP_TDC.1(VIRT).

7.1.8 Security Management

The security management facilities provided by the TOE are usable by authorized users and/or authorized administrators to modify the configuration of TSF. The configuration of TSF are hosted in the following locations:

- Configuration files (or TSF databases)
- Data structures maintained by the kernel and within the kernel memory

The TOE provides applications to authorized users as well as authorized administrators to perform various administrative tasks. These applications are documented as part of the administrator and user guidance. These applications are either used to modify configuration files or to access parameters controlled and enforced by the kernel via kernel-provided interfaces to user space.

Configuration options are stored in different configuration files. These files are protected using the DAC mechanisms against unauthorized access where usually the root user only is allowed to write to the files. In some special cases (like for `/etc/shadow`), the file is even readable to the root user only. It is the task of the persons responsible for setting up and administrating the system to ensure that the access control features of the TOE are used throughout the lifetime of the system to protect those databases. These configuration files are accessed using applications which are able to interpret the contents of these configuration files. Each TOE instance maintains its own TSF database. Synchronizing those databases is not performed in the evaluated configuration. If such synchronization is required by an organization it is the responsibility of an administrative user of the TOE to achieve this either manually or with some automated assistance.

To access data structures maintained by the kernel, applications use the kernel-provided interfaces, such as system calls, virtual file systems, netlink sockets, and device files. These kernel interfaces are restricted to authorized administrators or authorized users, if applicable, by either using DAC (for virtual file system objects) or special kernel-internal verification checks for each interface.

The TOE provides security management applications for all security-relevant settings listed throughout this ST, i.e. all FMT_MSA.1, FMT_MSA.3 and FMT_MTD.1 iterations, FMT_SMR.1, FMT_SMF.1.

7.1.8.1 Privileges

Privileges to perform administrative actions are maintained by the TOE. These privileges are separated into privileges to act on data or access functionality in user space and in kernel space.

Functionality accessible in user space are applications that can be invoked by users. Also, data accessible in user space is either data maintained with an application or data stored in persistent or transient storage objects. Privileges are controlled by permissions to invoke applications and to access data. For example, the configuration files including the user databases of `/etc/passwd` and `/etc/shadow` are accessible to the root user only. Therefore, the root user is given the privilege to perform modifications on this configuration data which constitutes administrative actions.

Functionality and data maintained by the kernel must be accessed using system calls. The kernel implements a privilege check for functions and data that shall not be accessible by normal users. These privileges are controlled with capabilities that can be assigned to processes. If a process is assigned with a capability, it is allowed to request special operations that other processes cannot. To implement consistency with the Unix legacy, processes with the effective UID of zero are implicitly given all capabilities. However, these processes may decide to drop capabilities. Such capabilities are marked by names with the prefix of "CAP_" throughout this document. The Linux kernel implements many more capabilities than mentioned in this document. These unmentioned capabilities protect functions that do not directly cover SFR functionality but need to be protected to ensure the integrity of the system and its resources.

8 Abbreviations, Terminology and References

8.1 Abbreviations

ACL

Access Control List

API

Application Programming Interface

KVM

Kernel Virtualized Machine

HTTP

Hypertext Transfer Protocol

SFR

Security Functional Requirement

SSL

Secure Sockets Layer

ST

Security Target

TCP/IP

Transmission Control Protocol / Internet Protocol

TLS

Transport Layer Security

TOE

Target of Evaluation

TSF

TOE Security Functionality

VM

Virtual Machine

VPN

Virtual Private Network

8.2 Terminology

This section contains definitions of technical terms that are used with a meaning specific to this document. Terms defined in the [CC] are not reiterated here, unless stated otherwise.

Authentication Data

Authentication data is the data used by users or remote entities to authenticate their claimed identity.

Authorized Administrator

This term refers to a user in one of the defined administrative roles of a Linux system. The TOE associates the user with the UID of zero and named "root" with administrative authorities. Effectively, the UID zero is assigned with all Linux capabilities known to the Linux kernel. Every user who is allowed to log on as that root user, or to switch their UID to the root user is considered an authorized administrator. In addition, any user who is able to execute applications which grant one or more Linux capabilities to be used in an unconditional manner is considered an authorized administrator. Note: the process executing on behalf of the root user must possess MLS override attributes to perform management aspects of the Authoritative Access Control Policy.

Category

A category is the non-hierarchical category of the lower MLS label defined with an SELinux label. Note: an SELinux label consists of four parts where the MLS label is one of them. The MLS label in turn is split into a higher and a lower MLS label part.

Classification

A sensitivity label associated with an object.

Clearance

A sensitivity label associated with a subject or user.

DAC

Discretionary Access Control implemented with permission bits and ACLs.

Data

Arbitrary bit sequences on persistent or transient storage media.

Dominates

Sensitivity label A dominates sensitivity label B if the hierarchical level of A is greater than or equal to the hierarchical level of B, and the category set of label A is a proper subset of or equal to the category set of label B. (cf. Incomparable sensitivity labels).

Guest

Software executing within a virtual machine environment. There can be zero or more guests executing concurrently on the host system.

Host

The host system provides the Linux environment that controls and manages the virtual machines. The host provides the execution environment for every virtual machine.

Information

Any data held within a server, including data in transit between systems.

IOMMU

Input / Output Memory Management Unit. This MMU allows the setup of multiple DMA areas for different virtual machines.

MLS

Multi-level security

Named Object

In Linux, those objects that are covered by access control policies. The list of objects defined as named objects is provided with FDP_ACC.1.

Object

For Linux, objects are defined by FDP_ACC.1.

OSPP

Operating System Protection Profile

OSPP EP

Operating System Protection Profile Extended Package

PAM

Pluggable Authentication Module - the authentication functionality provided with Linux is highly configurable by selecting and combining different modules implementing different aspects of the authentication process.

Product

The term product is used to define software components that comprise the Wind River Linux system.

QEMU

The QEMU software component implements the virtual devices and virtual resources for virtual machines. There is one instance of QEMU per virtual machine. The QEMU software component is also identified as the "kvm" application on the host system.

SELinux

Linux kernel LSM module that is able to implement arbitrary security policies. An SELinux policy distributed with the TOE implements multi-level or multi-category security.

Sensitivity Label

The TOE attaches a sensitivity label to each named object. This label consists of a hierarchical sensitivity level and a set of zero or more categories. The policy defines the number and names of the sensitivity levels and categories.

Subject

There are two classes of subjects in Red Hat Enterprise Linux: i) untrusted internal subject - this is a Linux process running on behalf of some user or providing an arbitrary service, running outside of the TSF (for example, with no privileges); ii) trusted internal subject - this is a Linux process running as part of the TSF (for example: service daemons and the process implementing the identification and authentication of users).

Target Of Evaluation (TOE)

The TOE is defined as the Red Hat Enterprise Linux operating system, running and tested on the hardware and firmware specified in this Security Target. The BootPROM firmware as well as the hardware are not part of the TOE.

User

Any individual/person or technical entity (such as a service added by the administrator on top of the TOE) who has a unique user identifier and who interacts with the Wind River Linux product.

User Security Attributes

Defined by functional requirement FIA_ATD.1, every user is associated with a number of security attributes which allow the TOE to enforce its security functions on this user. This also includes the user clearance which defines the maximum sensitivity label a user can have access to.

Virtual devices

See virtual resources for a generic explanation. This definition applies also to virtual devices, but with a focus to devices, such as disks, network cards, graphics cards, and similar.

Virtual machine

A virtual machine is an execution environment where the software executing within the virtual machine has access to the processor's user and supervisor state and resources defined by the host system. Resources include the number of processors, RAM size, physical devices, virtualized devices, communication channels to other virtual machines and the host system. For the KVM environment a virtual machine environment is controlled and provided by the Linux kernel hypervisor functionality plus the QEMU application instantiated for each virtual machine.

Virtual machine environment

See virtual machine.

Virtual resources

Virtual resources are resources that either do not physically exist and do not exist in the host system. Virtual resources are implemented by the virtual machine environment and are provided to the respective virtual machine. For example, virtual resources are special exceptions that can be triggered from the virtual machine environment to request services from the host system, such as para-virtualized drivers. Virtual devices can be considered one form of virtual resources.

8.3 References

CC	Common Criteria for Information Technology Security Evaluation
	Version 3.1R3
	Date July 2009
	Location http://www.commoncriteriaportal.org/files/ccfiles/CCPART1V3.1R3.pdf
	Location http://www.commoncriteriaportal.org/files/ccfiles/CCPART2V3.1R3.pdf
	Location http://www.commoncriteriaportal.org/files/ccfiles/CCPART3V3.1R3.pdf