

Red Hat Enterprise Linux 7 Virtualization Security Guide

Securing your virtual environment

Jiri HerrmannScott RadvanTahlia RichardsonThanks go to the following people for enabling the creation of this guide:Paul MooreKurt SeifriedDavid Jorm

Securing your virtual environment

Jiri Herrmann Red Hat Customer Content Services jherrman@redhat.com

Scott Radvan Red Hat Customer Content Services

Tahlia Richardson Red Hat Customer Content Services

Paul Moore Red Hat Engineering

Kurt Seifried Red Hat Engineering

David Jorm Red Hat Engineering

Thanks go to the following people for enabling the creation of this guide:

Legal Notice

Copyright © 2016 Red Hat, Inc.

This document is licensed by Red Hat under the <u>Creative Commons Attribution-ShareAlike 3.0</u> <u>Unported License</u>. If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ${\ensuremath{\mathbb R}}$ is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

 $MySQL \ \ensuremath{\mathbb{R}}$ is a registered trademark of $MySQL \ AB$ in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This guide provides an overview of virtualization security technologies provided by Red Hat. It also provides recommendations for securing hosts, guests, and shared infrastructure and resources in virtualized environments.

Table of Contents

Chapter 1. Introduction	2
1.1. Virtualized and Non-Virtualized Environments	2
1.2. Why Virtualization Security Matters	3
1.3. Leveraging SELinux with sVirt	4
Chapter 2. Host Security	5
2.1. Why Host Security Matters	5
2.2. Host Security Recommended Practices for Red Hat Enterprise Linux	5
Chapter 3. Guest Security	7
3.1. Why Guest Security Matters	7
3.2. Guest Security Recommended Practices	7
Chapter 4. sVirt	8
4.1. Introduction	8
4.2. SELinux and Mandatory Access Control (MAC)	8
4.3. sVirt Configuration	9
4.4. sVirt Labeling	10
Chapter 5. Network Security in a Virtualized Environment	13
5.1. Network Security Overview	13
5.2. Network Security Recommended Practices	13
Appendix A. Further Information	14
A.1. SELinux and sVirt	14
A.2. Virtualization Security	14
Appendix B. Revision History	15

Chapter 1. Introduction

1.1. Virtualized and Non-Virtualized Environments

A virtualized environment presents opportunities for both the discovery of new attack vectors and the refinement of existing exploits that may not previously have presented value to an attacker. It is therefore important to take steps to ensure the security of both the physical hosts and the guests running on them when creating and maintaining virtual machines.

Non-Virtualized Environment

In a non-virtualized environment, hosts are separated from each other physically and each host has a self-contained environment, consisting of services such as a web server, or a DNS server. These services communicate directly to their own user space, host kernel and physical host, offering their services directly to the network. The following image represents a non-virtualized environment:



Figure 1.1. Non-Virtualized Environment

Virtualized Environment

In a virtualized environment, several operating systems can be housed (as "guests") within a single host kernel and physical host. The following image represents a virtualized environment:



Figure 1.2. Virtualized Environment

When services are not virtualized, machines are physically separated. Any exploit is therefore usually contained to the affected machine, with the obvious exception of network attacks. When services are grouped together in a virtualized environment, extra vulnerabilities emerge in the system. If there is a security flaw in the hypervisor that can be exploited by a guest instance, this guest may be able to not only attack the host, but also other guests running on that host. This is not theoretical; attacks already exist on hypervisors. These attacks can extend beyond the guest instance and could expose other guests to attack.

1.2. Why Virtualization Security Matters

Deploying virtualization in your infrastructure provides many benefits but can also introduce new risks. Virtualized resources and services should be deployed with the following security considerations:

- The host/hypervisor become prime targets; they are often a single point of failure for guests and data.
- Virtual machines can interfere with each other in undesirable ways. Assuming no access controls were in place to help prevent this, one malicious guest could bypass a vulnerable hypervisor and directly access other resources on the host system, such as the storage of other guests.
- Resources and services can become difficult to track and maintain; with rapid deployment of virtualized systems comes an increased need for management of resources, including sufficient patching, monitoring and maintenance.
- Technical staff may lack knowledge, have gaps in skill sets, and have minimal experience in virtual environments. This is often a gateway to vulnerabilities.
- Resources such as storage can be spread across, and dependent upon, several machines. This can lead to overly complex environments, and poorly-managed and maintained systems.
- Virtualization does not remove any of the traditional security risks present in your environment; the entire solution stack, not just the virtualization layer, must be secured.

This guide aims to assist you in mitigating your security risks by offering a number of virtualization recommended practices for Red Hat Enterprise Linux that will help you secure your virtualized infrastructure.

1.3. Leveraging SELinux with sVirt

sVirt integrates virtualization into the existing security framework provided by SELinux (Security-Enhanced Linux), applying *Mandatory Access Control* (MAC) to virtual machines. The main objective of sVirt is to protect hosts and guests from attacks via security vulnerabilities in the hypervisor. SELinux secures a system by applying access policy across different processes. sVirt extends this capability to hosts and guests by treating each guest as a process, allowing administrators to apply similar policies designed to prevent malicious guests from accessing restricted resources. For more information on sVirt, refer to Chapter 4, sVirt.

Chapter 2. Host Security

2.1. Why Host Security Matters

When deploying virtualization technologies, the security of the host should be paramount. The Red Hat Enterprise Linux host system is responsible for managing and controlling access to the physical devices, storage and network as well as all virtualized guests themselves. If the host system is compromised, not only would the host system be vulnerable, but so would the guests and their data.

Virtualized guests are only as secure as their host system; securing the Red Hat Enterprise Linux host system is the first step towards ensuring a secure virtualization platform.

2.2. Host Security Recommended Practices for Red Hat Enterprise Linux

With host security being such a critical part of a secure virtualization infrastructure, the following recommended practices should serve as a starting point for securing a Red Hat Enterprise Linux host system:

- Run only the services necessary to support the use and management of your guest systems. If you need to provide additional services, such as file or print services, you should consider running those services on a Red Hat Enterprise Linux guest.
- Limit direct access to the system to only those users who have a need to manage the system. Consider disallowing shared root access and instead use tools such as **sudo** to grant privileged access to administrators based on their administrative roles.
- Ensure that SELinux is configured properly for your installation and is operating in enforcing mode. Besides being a good security practice, the advanced virtualization security functionality provided by sVirt relies on SELinux. Refer to <u>Chapter 4</u>, sVirt for more information on SELinux and sVirt.
- Ensure that auditing is enabled on the host system and that libvirt is configured to emit audit records. When auditing is enabled, libvirt will generate audit records for changes to guest configuration as well start/stop events which help you track the guest's state. In addition to the standard audit log inspection tools, the libvirt audit events can also be viewed using the specialized auvirt tool.
- Ensure that any remote management of the system takes place only over secured network channels. Tools such as SSH and network protocols such as TLS or SSL provide both authentication and data encryption to help ensure that only approved administrators can manage the system remotely.
- Ensure that the firewall is configured properly for your installation and is activated at boot. Only those network ports needed for the use and management of the system should be allowed.
- Refrain from granting guests direct access to entire disks or block devices (for example, /dev/sdb); instead, use partitions (for example, /dev/sdb1) or LVM volumes for guest storage.
- » Ensure that staff have adequate training and knowledge in virtual environments.



Warning

Attaching a USB device, Physical Function or physical device when SR-IOV is not available to a virtual machine could provide access to the device which is sufficient enough to overwrite that device's firmware. This presents a potential security issue by which an attacker could overwrite the device's firmware with malicious code and cause problems when moving the device between virtual machines or at host boot time. It is advised to use SR-IOV Virtual Function device assignment where applicable.

Note

The objective of this guide is to explain the unique security-related challenges, vulnerabilities, and solutions that are present in most virtualized environments, and the recommended method of addressing them. However, there are a number of recommended practices to follow when securing a Red Hat Enterprise Linux system that apply regardless of whether the system is a standalone, virtualization host, or guest instance. These recommended practices include procedures such as system updates, password security, encryption, and firewall configuration. This information is discussed in more detail in the *Red Hat Enterprise Linux Security Guide* which can be found at https://access.redhat.com/site/documentation/.

2.2.1. Special Considerations for Public Cloud Operators

Public cloud service providers are exposed to a number of security risks beyond that of the traditional virtualization user. Virtual guest isolation, both between the host and guest as well as between guests, is critical due to the threat of malicious guests and the requirements on customer data confidentiality and integrity across the virtualization infrastructure.

In addition to the Red Hat Enterprise Linux virtualization recommended practices previously listed, public cloud operators should also consider the following items:

- Disallow any direct hardware access from the guest. PCI, USB, FireWire, Thunderbolt, eSATA and other device passthrough mechanisms not only make management difficult, but often rely on the underlying hardware to enforce separation between the guests.
- Isolate the cloud operator's private management network from the customer guest network, and customer networks from one another, so that:
 - the guests cannot access the host systems over the network.
 - one customer cannot access another customer's guest systems directly via the cloud provider's internal network.

Chapter 3. Guest Security

3.1. Why Guest Security Matters

While the security of the host system is critical in ensuring the security of the guests running on the host, it does not remove the need for properly securing the individual guest machines. All of the security risks associated with a conventional, non-virtualized system still exist when the system is run as a virtualized guest. Any resources accessible to the guest system, such as critical business data or sensitive customer information, could be made vulnerable if the guest system were to be compromised.

3.2. Guest Security Recommended Practices

All of the recommended practices for securing a Red Hat Enterprise Linux system documented in the *Red Hat Enterprise Linux Security Guide* apply to conventional, non-virtualized systems as well as systems installed as a virtualized guest. However, there are a few security practices which are of critical importance when running guests in a virtualized environment:

- With all management of the guest likely taking place remotely, ensure that the management of the system takes place only over secured network channels. Tools such as SSH and network protocols such as TLS or SSL provide both authentication and data encryption to ensure that only approved administrators can manage the system remotely.
- Some virtualization technologies use special guest agents or drivers to enable some virtualization specific features. Ensure that these agents and applications are secured using the standard Red Hat Enterprise Linux security features, such as SELinux.
- In virtualized environments there is a greater risk of sensitive data being accessed outside the protection boundaries of the guest system. Protect stored sensitive data using encryption tools such as dm-crypt and GnuPG; although special care needs to be taken to ensure the confidentiality of the encryption keys.

Note

Using page deduplication technology such as Kernel Same-page Merging (KSM) may introduce side channels that could potentially be used to leak information across guests. In situations where this is a concern, KSM can be disabled either globally or on a per-guest basis. See the <u>Red Hat Enterprise Linux 7 Virtualization Tuning and Optimization Guide</u> for more information about KSM.

Chapter 4. sVirt

4.1. Introduction

Since virtual machines under KVM are implemented as Linux processes, KVM leverages the standard Linux security model to provide isolation and resource controls. The Linux kernel includes SELinux (Security-Enhanced Linux), a project developed by the US National Security Agency to add mandatory access control (MAC), multi-level security (MLS) and multi-category security (MCS) through a flexible and customizable security policy. SELinux provides strict resource isolation and confinement for processes running on top of the Linux kernel, including virtual machine processes. The sVirt project builds upon SELinux to further facilitate virtual machine isolation and controlled sharing. For example, fine-grained permissions could be applied to group virtual machines together to share resources.

From a security point of view, the hypervisor is a tempting target for attackers, as a compromised hypervisor could lead to the compromise of all virtual machines running on the host system. Integrating SELinux into virtualization technologies helps improve hypervisor security against malicious virtual machines trying to gain access to the host system or other virtual machines.

Refer to the following image which represents isolated guests, limiting the ability for a compromised hypervisor (or guest) to launch further attacks, or to extend to another instance:



Figure 4.1. Attack path isolated by SELinux



4.2. SELinux and Mandatory Access Control (MAC)

Security-Enhanced Linux (SELinux) is an implementation of MAC in the Linux kernel, checking for

allowed operations after standard discretionary access controls (DAC) are checked. SELinux can enforce a user-customizable security policy on running processes and their actions, including attempts to access file system objects. Enabled by default in Red Hat Enterprise Linux, SELinux limits the scope of potential damage that can result from the exploitation of vulnerabilities in applications and system services, such as the hypervisor.

sVirt integrates with libvirt, a virtualization management abstraction layer, to provide a MAC framework for virtual machines. This architecture allows all virtualization platforms supported by libvirt and all MAC implementations supported by sVirt to interoperate.

4.3. sVirt Configuration

SELinux Booleans are variables that can be toggled on or off, quickly enabling or disabling features or other special conditions. Booleans can be toggled by running either **setsebool boolean_name** {**on**|**off**} for a temporary change, or **setsebool** -**P boolean_name** {**on**|**off**} to make the change persistent across reboots.

The following table shows the SELinux Boolean values that affect KVM when launched by libvirt. The current state of these booleans (on or off) can be found by running the command **getsebool** - **a|grep virt**.

SELinux Boolean	Description
staff_use_svirt	Allow staff user to create and transition to sVirt domains.
unprivuser_use_svirt	Allow unprivileged user to create and transition to sVirt domains.
virt_sandbox_use_audit	Allow sandbox containers to send audit messages.
virt_sandbox_use_netlink	Allow sandbox containers to use netlink system calls.
virt_sandbox_use_sys_admin	Allow sandbox containers to use sys_admin system calls, such as mount.
virt_transition_userdomain	Allow virtual processes to run as user domains.
virt_use_comm	Allow virt to use serial/parallel communication ports.
virt_use_execmem	Allow confined virtual guests to use executable memory and executable stack.
virt_use_fusefs	Allow virt to read FUSE mounted files.
virt_use_nfs	Allow virt to manage NFS mounted files.
virt_use_rawip	Allow virt to interact with rawip sockets.
virt_use_samba	Allow virt to manage CIFS mounted files.
virt_use_sanlock	Allow confined virtual guests to interact with the sanlock.
virt_use_usb	Allow virt to use USB devices.
virt_use_xserver	Allow virtual machine to interact with the X Window System.

Table 4.1. KVM SELinux Booleans

Note

For more information on SELinux Booleans, refer to the *Red Hat Enterprise Linux SELinux Users* and *Administrators Guide* at https://access.redhat.com/site/documentation/.

4.4. sVirt Labeling

Like other services under the protection of SELinux, sVirt uses process based mechanisms, labels and restrictions to provide extra security and control over guest instances. Labels are applied automatically to resources on the system based on the currently running virtual machines (dynamic), but can also be manually specified by the administrator (static), to meet any specific requirements that may exist.

4.4.1. Types of sVirt Labels

The following table outlines the different sVirt labels that can be assigned to resources such as virtual machine processes, image files and shared content:

Туре	SELinux Context	Description/Effect
Virtual Machine Processes	system_u:system_r:svirt_t:MCS1	<i>MCS1</i> is a randomly selected field. Currently approximately 500,000 labels are supported.
Virtual Machine Image	system_u:object_r:svirt_image_t : <i>MCS1</i>	Only <i>svirt_t</i> processes with the same <i>MCS1</i> fields are able to read/write these image files and devices.
Virtual Machine Shared Read/Write Content	system_u:object_r:svirt_image_t :s0	All <i>svirt_t</i> processes are allowed to write to the <i>svirt_image_t:s0</i> files and devices.
Virtual Machine Shared Shared Read Only content	system_u:object_r:svirt_content _t:s0	All <i>svirt_t</i> processes are able to read files/devices with this label.
Virtual Machine Image	system_u:object_r:virt_content_t :s0	System default label used when an image exits. No <i>svirt_t</i> virtual processes are allowed to read files/devices with this label.

Table 4.2. sVirt Labels

4.4.2. Dynamic Configuration

Dynamic label configuration is the default labeling option when using sVirt with SELinux. Refer to the following example which demonstrates dynamic labeling:

```
# ps -eZ | grep qemu-kvm
```

system_u:system_r:svirt_t:s0:c87,c520 27950 ? 00:00:17 qemu-kvm

In this example, the **qemu-kvm** process has a base label of **system_u: system_r: svirt_t: s0**. The libvirt system has generated a unique MCS label of **c87**, **c520** for this process. The base label and the MCS label are combined to form the complete security label for the process. Likewise, libvirt takes the same MCS label and base label to form the image label. This image label is then automatically applied to all host files that the VM is required to access, such as disk images, disk devices, PCI devices, USB devices, and kernel/initrd files. Each process is isolated from other virtual machines with different labels.

The following example shows the virtual machine's unique security label (with a corresponding MCS label of **c87**, **c520** in this case) as applied to the guest disk image file in /var/lib/libvirt/images:

```
# ls -lZ /var/lib/libvirt/images/*
system_u:object_r:svirt_image_t:s0:c87,c520 image1
```

The following example shows dynamic labeling in the XML configuration for the guest:

```
<seclabel type='dynamic' model='selinux' relabel='yes'>
    <label>system_u:system_r:svirt_t:s0:c87,c520</label>
    <imagelabel>system_u:object_r:svirt_image_t:s0:c87,c520</imagelabel>
</seclabel>
```

4.4.3. Dynamic Configuration with Base Labeling

To override the default base security label in dynamic mode, the **<baselabel>** option can be configured manually in the XML guest configuration, as shown in this example:

```
<seclabel type='dynamic' model='selinux' relabel='yes'>
    <baselabel>system_u:system_r:svirt_custom_t:s0</baselabel>
    <label>system_u:system_r:svirt_custom_t:s0:c87,c520</label>
    <imagelabel>system_u:object_r:svirt_image_t:s0:c87,c520</imagelabel>
</seclabel>
```

4.4.4. Static Configuration with Dynamic Resource Labeling

Some applications require full control over the generation of security labels but still require libvirt to take care of resource labeling. The following guest XML configuration demonstrates an example of static configuration with dynamic resource labeling:

4.4.5. Static Configuration without Resource Labeling

Primarily used in MLS (multi-level security) or otherwise strictly controlled environments, static configuration without resource relabeling is possible. Static labels allow the administrator to select a specific label, including the MCS/MLS field, for a virtual machine. Administrators who run statically-labeled virtual machines are responsible for setting the correct label on the image files. The virtual machine will always be started with that label, and the sVirt system will never modify the label of a statically-labelled virtual machine's content. The following guest XML configuration demonstrates an example of this scenario:

```
<seclabel type='static' model='selinux' relabel='no'>
    <label>system_u:system_r:svirt_custom_t:s0:c87,c520</label>
</seclabel>
```

Chapter 5. Network Security in a Virtualized Environment

5.1. Network Security Overview

In almost all situations, the network is the only way to access systems, applications, and management interfaces. As networking plays such a critical role in the management of virtualized systems and the availability of their hosted applications, it is very important to ensure that the network channels both to and from the virtualized systems are secure.

Securing the network allows administrators to control access and protect sensitive data from information leaks and tampering.

5.2. Network Security Recommended Practices

Network security is a critical part of a secure virtualization infrastructure. Refer to the following recommended practices for securing the network:

- Ensure that remote management of the system takes place only over secured network channels. Tools such as SSH and network protocols such as TLS or SSL provide both authentication and data encryption to assist with secure and controlled access to systems.
- Ensure that guest applications transferring sensitive data do so over secured network channels. If protocols such as TLS or SSL are not available, consider using one like IPsec.
- Configure firewalls and ensure they are activated at boot. Only network ports needed for the use and management of the system should be allowed. Test and review firewall rules regularly.

5.2.1. Securing Connectivity to SPICE

The SPICE remote desktop protocol supports SSL/TLS which should be enabled for all of the SPICE communication channels (main, display, inputs, cursor, playback, record).

5.2.2. Securing Connectivity to Storage

You can connect virtualized systems to networked storage in many different ways. Each approach presents different security benefits and concerns, however the same security principles apply to each: authenticate the remote store pool before use, and protect the confidentiality and integrity of the data while it is being transferred.

The data must also remain secure while it is stored. Red Hat recommends that data is encrypted or digitally signed before storing, or both.

Note

For more information on networked storage, refer to the Storage Pools chapter of the <u>Red Hat</u> <u>Enterprise Linux Virtualization Deployment and Administration Guide</u>.

Appendix A. Further Information

A.1. SELinux and sVirt

Further information on SELinux and sVirt:

- Main SELinux website: http://www.nsa.gov/research/selinux/index.shtml.
- SELinux documentation: http://www.nsa.gov/research/selinux/docs.shtml.
- Main sVirt website: http://selinuxproject.org/page/SVirt.
- » Dan Walsh's blog: http://danwalsh.livejournal.com/.
- The unofficial SELinux FAQ: http://www.crypt.gen.nz/selinux/faq.html.

A.2. Virtualization Security

Further information on virtualization security:

NIST (National Institute of Standards and Technology) full virtualization security guidelines: http://www.nist.gov/itl/csd/virtual-020111.cfm.

Appendix B. Revision History

Revision 1.0-15	Mon Oct 17 2016	Jiri Herrmann
Version for 7.3 GA publication		
Revision 1.0-9	Thu Oct 08 2015	Jiri Herrmann
Cleaned up the Revision Histor	у	
Revision 1.0-8	Wed Feb 18 2015	Scott Radvan
Version for 7.1 GA release.		