Red Hat Enterprise Linux Release 9.2 Manual Pages on 'kernel_lockdown.7' command

$ man kernel_lockdown.7

KERNEL_LOCKDOWN(7)   Linux Programmer's Manual   KERNEL_LOCKDOWN(7)

NAME

   kernel_lockdown - kernel image access prevention feature

DESCRIPTION

   The Kernel Lockdown feature is designed to prevent both direct and in?

   direct access to a running kernel image, attempting to protect against

   unauthorized modification of the kernel image and to prevent access to

   security and cryptographic data located in kernel memory, whilst still

   permitting driver modules to be loaded.

   Lockdown is typically enabled during boot and may be terminated, if

   configured, by typing a special key combination on a directly attached

   physical keyboard.

   If a prohibited or restricted feature is accessed or used, the kernel

   will emit a message that looks like:

        Lockdown: X: Y is restricted, see man kernel_lockdown.7

   where X indicates the process name and Y indicates what is restricted.

   On an EFI-enabled x86 or arm64 machine, lockdown will be automatically

   enabled if the system boots in EFI Secure Boot mode.

   If the kernel is appropriately configured, lockdown may be lifted by

   typing the appropriate sequence on a directly attached physical key?

   board. For x86 machines, this is SysRq+x.

 Coverage

   When lockdown is in effect, a number of features are disabled or have

their use restricted. This includes special device files and kernel services that allow direct access of the kernel image:

    /dev/mem

    /dev/kmem

    /dev/kcore

    /dev/ioports

    BPF

    kprobes

and the ability to directly configure and control devices, so as to prevent the use of a device to access or modify a kernel image:

? The use of module parameters that directly specify hardware parame?
  ters to drivers through the kernel command line or when loading a
  module.

? The use of direct PCI BAR access.

? The use of the ioperm and iopl instructions on x86.

? The use of the KD*IO console ioctls.

? The use of the TIOCSSERIAL serial ioctl.

? The alteration of MSR registers on x86.

? The replacement of the PCMCIA CIS.

? The overriding of ACPI tables.

? The use of ACPI error injection.

? The specification of the ACPI RDSP address.

? The use of ACPI custom methods.

Certain facilities are restricted:

? Only validly signed modules may be loaded (waived if the module file
  being loaded is vouched for by IMA appraisal).

? Only validly signed binaries may be kexec'd (waived if the binary im?
  age file to be executed is vouched for by IMA appraisal).

? Unencrypted hibernation/suspend to swap are disallowed as the kernel
  image is saved to a medium that can then be accessed.

? Use of debugfs is not permitted as this allows a whole range of ac?
  tions including direct configuration of, access to and driving of
  hardware.

? IMA requires the addition of the "secure_boot" rules to the policy,

  whether or not they are specified on the command line, for both the

  built-in and custom policies in secure boot lockdown mode.

VERSIONS

The Kernel Lockdown feature was added in Linux 5.4.

NOTES

The Kernel Lockdown feature is enabled by CONFIG_SECURITY_LOCKDOWN_LSM.

The lsm=lsm1,...,lsmN command line parameter controls the sequence of

the initialization of Linux Security Modules. It must contain the

string lockdown to enable the Kernel Lockdown feature. If the command

line parameter is not specified, the initialization falls back to the

value of the deprecated security= command line parameter and further to

the value of CONFIG_LSM.

COLOPHON

This page is part of release 5.10 of the Linux man-pages project. A

description of the project, information about reporting bugs, and the

latest   version   of   this   page,   can   be   found   at

https://www.kernel.org/doc/man-pages/.

Linux   2020-11-01   KERNEL_LOCKDOWN(7)