


Linux Security and Isolation APIs Essentials

Course code: M7D-SISESS01

This course provides an overview of the low-level Linux features—capabilities, namespaces, and control groups (cgroups)—that are used to build containers and sandboxes. Beginning with classical privileged programs (set-UID-*root* programs), we look at how capabilities and namespaces can be used to place processes in “a world of their own” in which they have private instances of various “global” resources. Those features—user namespaces in particular—can also be used to implement the notion of a process that is superuser inside a container while being unprivileged outside the container. Finally, we’ll see how cgroups can be used to limit resource consumption, so that the processes in a container can’t negatively impact other users on the system.



Audience and prerequisites

The primary audience comprises designers and programmers building privileged applications, container applications, and sandboxing applications. Systems administrators who manage such applications will also find the course of benefit.

Participants should have some awareness of classical Linux/UNIX concepts such as file descriptors and file I/O, signals, and the process lifecycle (*fork()*, *exec()*, *wait()*, *exit()*). In addition, participants should have a reading knowledge of the C programming language. (Note, however, that the course exercises do not require writing any programs.)

Related courses

The *Linux Security and Isolation APIs* (M7D-SECISOL02) course covers the same topics as this course, but in greater depth.

Course materials

- Course books (written by the trainer) that include all slides and exercises presented in the course

- An electronic copy of the trainer’s book, *The Linux Programming Interface*
- Numerous example programs written by the course trainer

Course duration and format

One day, with around 20% of the course time devoted to practical sessions.

Course inquiries and bookings

For inquiries about courses and consulting, you can contact us in the following ways:

- Email: training@man7.org
- Phone: +49 (89) 2155 2990 (German landline)

Prices, dates, and further details

For course prices, upcoming course dates, and further information about the course, please visit the course web page, <http://man7.org/training/secisolintro/>.

About the trainer



Michael Kerrisk has a unique set of qualifications and experience that ensure that course participants receive training of a very high standard:

- He has been programming on UNIX systems since 1987 and began teaching UNIX system programming courses in 1989.
- He is the author of *The Linux Programming Interface*, a 1550-page book acclaimed as the definitive work on Linux system programming.

- He has been actively involved in Linux development, working with kernel developers on testing, review, and design of new Linux kernel–user-space APIs.
- Since 2000, he has been involved in the Linux *man-pages* project, which provides the manual pages documenting Linux system calls and C library APIs, and was the project maintainer from 2004 to 2021.

Linux Security and Isolation APIs Essentials: course contents in detail

Topics marked with an asterisk (*) may be covered, if time permits.

1. Course Introduction

2. Classical Privileged Programs

- A simple set-user-ID program
- Saved set-user-ID and saved set-group-ID
- Changing process credentials
- A few guidelines for writing privileged programs

3. Capabilities

- Process and file capabilities
- Permitted and effective capabilities
- Setting and viewing file capabilities
- Text-form capabilities
- Capabilities and `execve()`
- Capabilities and UID transitions

4. Namespaces

- An example: UTS namespaces
- Namespaces commands
- Namespaces demonstration (UTS namespaces)
- Namespace types and APIs
- Mount namespaces
- PID namespaces

5. Namespaces APIs (*)

- API Overview

- Creating a child process in new namespaces: `clone()`

6. User Namespaces

- Overview of user namespaces
- Creating and joining a user namespace
- User namespaces: UID and GID mappings
- Accessing files (and other objects with UIDs/GIDs)
- Combining user namespaces with other namespaces

7. User Namespaces and Capabilities

- User namespaces and capabilities
- What does it mean to be superuser in a namespace?

8. Cgroups: Introduction

- Preamble
- What are control groups?
- An example: the `pids` controller
- Creating and destroying cgroups
- Populating a cgroup
- Enabling and disabling controllers

9. Cgroups: Other Controllers (*)

- The `cpu` and `freezer` controllers

The following are some of the **other courses taught by Michael Kerrisk**. Custom courses are also available upon request. Further details on these and other courses can be found at <http://man7.org/training/>. For course inquiries please email training@man7.org or phone +49 (89) 2155 2990 (German landline).

Linux Security and Isolation APIs

Course code: M7D-SECISOL02 (4 days)

Covering topics including control cgroups (cgroups), namespaces (with a deep dive into user namespaces), capabilities, and seccomp (secure computing), this course provides a deep understanding of the low-level Linux features used to design, build, and troubleshoot container, virtualization, and sandboxing frameworks. [This course is an expanded version of the course described above.]

Linux/UNIX Network Programming

Course code: M7D-NWP03 (3 days)

This course covers sockets programming (both UNIX and Internet domain sockets), and the use of relevant I/O techniques for working with sockets (`poll()`, `epoll`, nonblocking I/O). In addition, we look at the TCP/IP protocol stack (including details of TCP such as the 3-way handshake and the TCP state machine), the use of monitoring and tracing tools (`ss`, `netstat`, and `tcpdump/wireshark`), and raw sockets.

Linux/UNIX System Programming

Course code: M7D-LUSP01 (5 days)

This course covers the APIs used to build system-level applications on Linux and UNIX systems ranging from embedded processors to enterprise servers. The presentations and practical exercises provide participants with the knowledge needed to write complex system, network, and multithreaded applications. Topics include: file I/O; signals; process creation and termination; program execution; POSIX threads; interprocess communication, and I/O multiplexing (`poll()`, `epoll`).

Building and Using Shared Libraries on Linux

Course code: M7D-SHLIB04 (2.5 days)

This course describes how to design, build, and use shared libraries on Linux. Topics include: fundamentals of library creation and use; shared library versioning; symbol resolution; library search order; executable and linking format (ELF); dynamically loaded libraries; controlling symbol visibility; and symbol versioning.