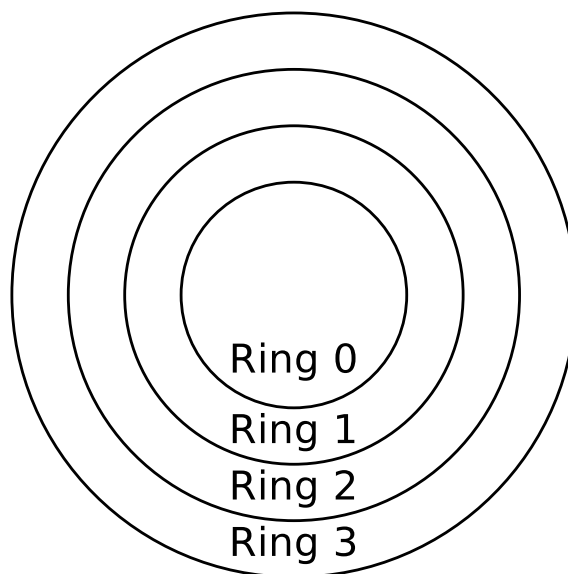


Exercise Sheet 7

Exercise 1 (System Calls)

1. x86-CPUs contain 4 privilege levels („rings“) for processes. Mark in the diagram (*clearly visible!*) the kernel mode and the user mode.




2. Name the ring to which the operating system is assigned. Name the ring to which the user applications are assigned.
3. Name the ring to which processes are assigned that have full access to the hardware.
4. Name a reason for the differentiation between user mode and kernel mode.
5. Explain what a system call is.
6. Explain what a context switch is.
7. Name two reasons why user mode processes should not call system calls directly.
8. What alternatives exist, if user mode processes should not call system calls directly?

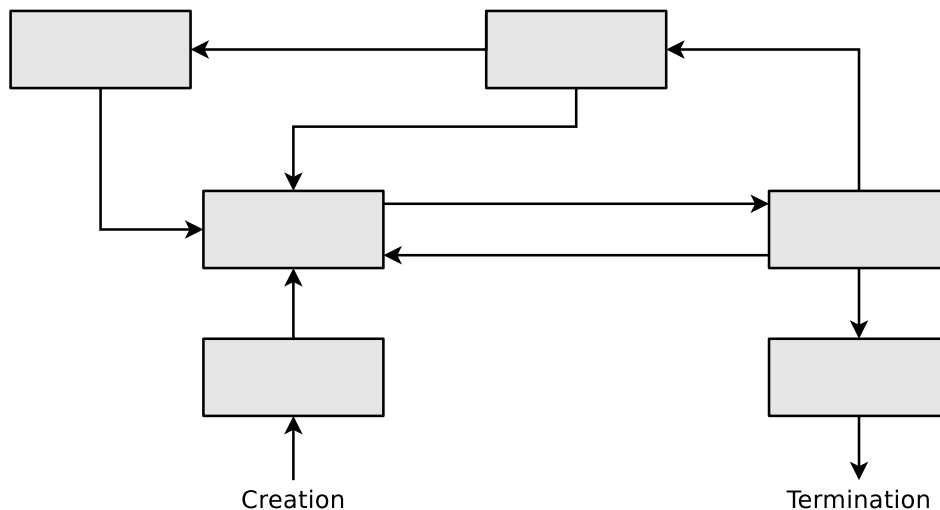
Exercise 2 (Processes)

1. Name the three sorts of process context information the operating system stores.

2. Name the process context information that are not stored in the process control block.
3. Explain why the process control block does not store all process context information.
4. Explain the task of the dispatcher.
5. Explain the task of the scheduler.
6. The process state model with 2 states is the smallest possible process model. Enter the names of the states in the diagram of the process state model with 2 states.

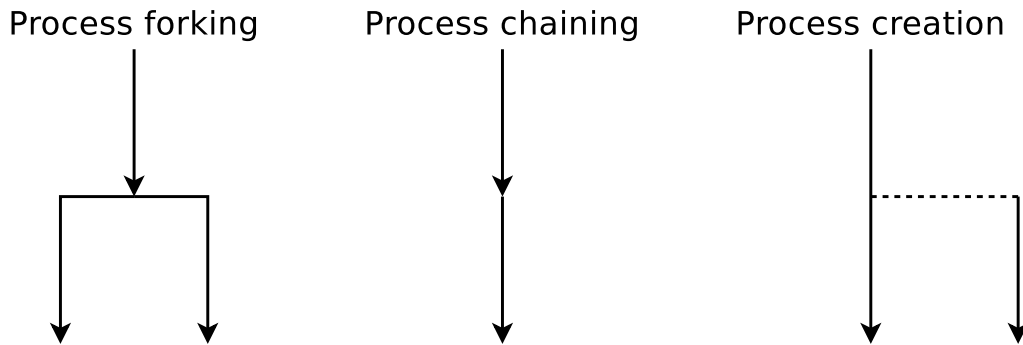


7. Does the process state model with 2 states make sense? Explain your answer shortly.
8. Enter the names of the states in the diagram of the process state model with 6 states.



9. Explain what a zombie process is.
10. Explain the task of the process table.
11. Give the number of status lists, the operating system manages for processes in „blocked“ state.
12. Explain what happens if a new process is to be created, but the operating system has no more free process IDs (PID) left.

13. Describe the effect of calling the system call `fork()`.
14. Describe the effect of calling the system call `exec()`.
15. The three diagrams below show all existing ways of creating a new process. Specify for each diagram, which system call(s) are required to implement the illustrated way of process creation.



16. A parent process (PID = 75) with the characteristics, described in the table below, creates a child process (PID = 198) by using the system call `fork()`. Enter the four missing values into the table.

	Parent Process	Child Process
PPID	72	
PID	75	198
UID	18	
Return value of <code>fork()</code>		

17. Describe what `init` is and what its task is.
18. Name the differences of a child process from the parent process shortly after its creation.
19. Describe the effect, when a parent process is terminated before the child process.
20. Describe what data the Text Segment contains.
21. Describe what data the Heap contains.
22. Describe what data the Stack contains.

Exercise 3 (Time-based Command Execution, Control Structures, Archiving)

1. Program a shell script, which reads two numbers as command line arguments. The script should check whether the numbers are identical, and print out the result of the check.
2. Extend the shell script in a way that if the numbers are not identical, it is checked, which one of the two numbers is the larger one. The result of the check should be printed out.
3. Program a shell script, which creates a backup of a directory of your choice. The script should create an archive file with the file extension `.tar.bz2` from the directory. The archive file should be stored in the directory `/tmp`. The name of the archive file should correspond to the following naming scheme:

`Backup_<USERNAME>_<YEAR>_<MONTH>_<DAY>.tar.bz2`

The fields `<USERNAME>`, `<YEAR>`, `<MONTH>` and `<DAY>` should be replaced by the current values.

4. Program a shell script, which checks if already today an archive file was created according to the naming scheme of subtask 3. The result of the check should be printed out in the shell.
5. Write two `cron` jobs. The first `cron` job should execute at 6:15 am on every day (except on weekends) the shell script from subtask 3, which creates the archive file with the backup.

The second `cron` job should execute at 11:45 am on every day (except on weekends) the shell script from subtask 4, which checks, whether already today an archive file was created.

The output from the shell scripts should be appended to a file `/tmp/Backup-Log.txt`. If the archive file `Backup...tar.bz2` has been created successfully, this should be noted in the log file `/tmp/Backup-Log.txt`.

Before each new entry in the file, lines according to the following pattern (with current values) should be inserted into the log file `/tmp/Backup-Log.txt`.

```
*****  
20.11.2013 --- Time: 21:39:51
```

Exercise 4 (Shell Scripts)

1. Program a shell script, which checks for a file, which is specified as an argument, whether it exists and if it is a file, a directory, a symbolic link, a socket or a named pipe.
 - The script should print out the result of the check.
2. Extend the shell script from subtask 1 in a way that if the file, which is specified as an argument, exists, it is checked, if the file could be executed and if write access would be possible.
3. Program a shell script, which reads so long text on the command line, until it is terminated by typing **END**.
 - The script should convert the text, which is read in from the command line, to uppercase.
4. Program a shell script, which prints out the number of running processes for all logged in users.
5. Extend the shell script from subtask 4 in a way that that the output is sorted.
 - The user with most processes should stand at the beginning.
6. Program a shell script, which checks after start every 10 seconds, if a file `/tmp/lock.txt` exists.
 - Each time after the script has checked the existence of the file, it should output an appropriate message on the shell.
 - Once the file `/tmp/lock.txt` exists, the script should terminate itself.