

ECE 5984 Virtualization Technologies

# Intermission

Pierre Olivier

# Outline

- 1) Textbook link
- 2) Unanswered questions
- 3) Project discussion session

# Outline

- 1) Textbook link
- 2) Unanswered questions
- 3) Project discussion session

# Textbook link

- 1) Go to <https://lib.vt.edu/>
- 2) Search for the textbook title
- 3) First link points to a webpage where you can get the PDF

# Outline

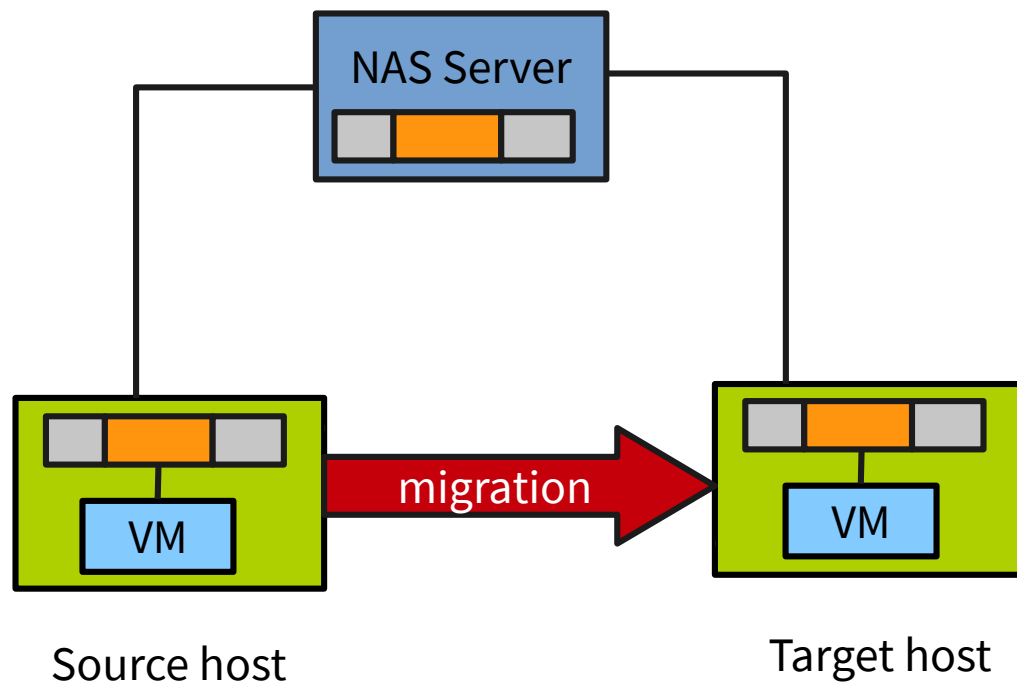
- 1) Textbook link
- 2) Unanswered questions
- 3) Project discussion session

# Unanswered Questions

## I/O and VM migration

### ■ Storage:

- ◆ Use Network-Attached-Storage (NAS) on the hosts



# Unanswered Questions

## I/O and VM migration

### ■ Network:


- ◆ Migrated VM includes connections state and will keep its IP
- ◆ Assume source and target host exist on the same switched subnet
- ◆ Send an ARP reply from the source host indicating that the IP moved to a new physical location
  - A few packets are lost in the process, but network protocols are resilient against that (packets will be resent transparently)

<http://www.cl.cam.ac.uk/research/srg/netos/papers/2005-migration-nsdi-pre.pdf>

# Unanswered Questions

## NUMA Abstraction

- Hardware provides cache-coherent access to all memory from each processor
- OS generally bootstraps on a single core (ex socket 0 core 0) then initializes all other cores from all other sockets
- At boot time the OS determines the topology:
  - ◆ Logical CPU/Core/Chip/NUMA domain + Amount of memory for each NUMA domain
  - ◆ Allows to take intelligent decisions for memory allocation





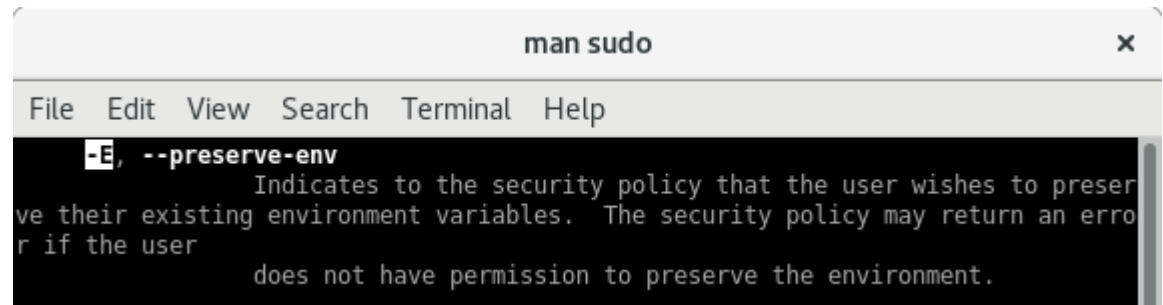
# Outline

- 1) Textbook link
- 2) Unanswered questions
- 3) **Project discussion session**

# Project discussion session

## ■ A few mistakes have been corrected in the technical guide

- ◆ C wrapper for *rename* syscall
- ◆ User environment variables are not carried out with sudo!
  - **VAR=VAL sudo <command>**
    - **command** does not have access to user-defined **VAR**
  - Solution: **sudo VAR=VAL <command>**
  - Another solution:
    - **export VAL=VAR**
    - **sudo -E <command>**



```
man sudo
File Edit View Search Terminal Help
-E, --preserve-env
    Indicates to the security policy that the user wishes to preserve their existing environment variables. The security policy may return an error if the user does not have permission to preserve the environment.
```

# Project discussion session

- Some students are still in search for a partner
  - ◆ 3 (remote) students contacted me by now
  - ◆ Any other contact me before tomorrow!
- Register your group!
  - ◆ <https://canvas.vt.edu/courses/63785/groups>

# Design discussion

# Project design discussion

## Postmark

### ■ Porting postmark

- ◆ Mkdir
- ◆ Rmdir
- ◆ Getwd

### ■ Other syscalls

- ◆ Remove
  - Not a syscall
  - Unlink?

```
pierre@orchid: ~/Desktop/HermitCore/build
File Edit View Search Terminal Tabs Help
pierre@orchid: ~/Desktop/HermitCo... x pierre@orchid: ~/Desktop/HermitCo... x
[ 8%] Built target hellof
[ 16%] Built target hello++
[ 25%] Built target hello
[ 33%] Built target server
[ 41%] Built target jacobi
Scanning dependencies of target postmark
[ 45%] Building C object CMakeFiles/postmark.dir/postmark.c.obj
[ 50%] Linking C executable postmark
CMakeFiles/postmark.dir/postmark.c.obj: In function `create_subdirectories.part.9':
postmark.c:(.text+0x588): undefined reference to `mkdir'
CMakeFiles/postmark.dir/postmark.c.obj: In function `delete_subdirectories.part.10':
postmark.c:(.text+0x613): undefined reference to `rmdir'
CMakeFiles/postmark.dir/postmark.c.obj: In function `cli_show':
postmark.c:(.text+0x10a4): undefined reference to `getwd'
collect2: error: ld returned 1 exit status
CMakeFiles/postmark.dir/build.make:94: recipe for target 'postmark' failed
make[5]: *** [postmark] Error 1
CMakeFiles/Makefile2:253: recipe for target 'CMakeFiles/postmark.dir/all' failed
make[4]: *** [CMakeFiles/postmark.dir/all] Error 2
Makefile:129: recipe for target 'all' failed
make[3]: *** [all] Error 2
CMakeFiles/tests.dir/build.make:112: recipe for target 'tests-prefix/src/tests-stamp/tests-build' failed
make[2]: *** [tests-prefix/src/tests-stamp/tests-build] Error 2
CMakeFiles/Makefile2:665: recipe for target 'CMakeFiles/tests.dir/all' failed
make[1]: *** [CMakeFiles/tests.dir/all] Error 2
Makefile:151: recipe for target 'all' failed
make: *** [all] Error 2
pierre@orchid ~/Desktop/HermitCore/build <master*>
$
```

# Project design discussion

## ■ Other syscalls?

- ◆ One idea is to strace a regular Linux postmark run

```
pierre@orchid: ~/Desktop/HermitCore/usr/tests
File Edit View Search Terminal Help
└─$ gcc postmark.c -o postmark
postmark.c: In function 'cli_show':
postmark.c:1109:7: warning: 'getwd' is deprecated [-Wdeprecated-declarations]
      (file_system_count==0)?GETWD(current_dir):");
      ^
In file included from postmark.c:58:0:
/usr/include/unistd.h:528:14: note: declared here
extern char *getwd (char *__buf)
              ^
└─$ strace -c ./postmark
PostMark v1.51 : 8/14/01
pm>set number 10
pm>set transactions 10
pm>run
Creating files...Done
Performing transactions.....Done
Deleting files...Done
Time:
    1 seconds total
    1 seconds of transactions (10 per second)
Files:
    16 created (16 per second)
        Creation alone: 10 files (10 per second)
        Mixed with transactions: 6 files (6 per second)
    5 read (5 per second)
    5 appended (5 per second)
    16 deleted (16 per second)
        Deletion alone: 12 files (12 per second)
        Mixed with transactions: 4 files (4 per second)
Data:
    35.59 kilobytes read (35.59 kilobytes per second)
    87.58 kilobytes written (87.58 kilobytes per second)
pm>% time      seconds  usecs/call   calls  errors syscall
-----
58.16  0.000435      27      16      0      unlink
37.57  0.000281       4      71      0      write
 4.28  0.000032       2      17      0      read
 0.00  0.000000       0      28      0      close
 0.00  0.000000       0      30      0      fstat
 0.00  0.000000       0       5      0      lseek
 0.00  0.000000       0       5      0      mmap
 0.00  0.000000       0       4      0      mprotect
 0.00  0.000000       0       1      0      munmap
 0.00  0.000000       0       3      0      brk
 0.00  0.000000       0       3      3      access
 0.00  0.000000       0       1      0      execve
 0.00  0.000000       0       1      0      arch_prctl
 0.00  0.000000       0      29      1      openat
-----
100.00  0.000748      214     214      4      total
└─$
```