Linux is obsolete 2.0 

Rüdiger Weis

TFH Berlin

CCCamp 2007

Tanenbaum versus Brown

''Thus, of course, Linus didn't sit down in a vacuum and suddenly type in the Linux source code. He had my book, was running MINIX, and undoubtedly knew the history (since it is in my book). But the code was his.
The proof of this is that he messed the design up.''

http://www.cs.vu.nl/~ast/brown/

USENIX ;login:

USENIX April 2006, Rick Farrow, "Musings"

`http://www.usenix.com/publications/login/2006-04/openpdfs/musings.pdf`

> *"While I have been busy ranting about the need for new operating system design, Andrew Tanenbaum and his students have been busy writing MINIX."*

Minix3

`http://www.minix3.org/`

MINIX 3 is initially targeted at the following areas:

- Applications where very high reliability is required
- Single-chip, small-RAM, low-power, $100 laptops
- Embedded systems
- Education (e.g., operating systems courses at universities)

Minix3 Features

- POSIX compliant
- Full C source code supplied under a BSD-type licence.
- Networking with TCP/IP
- X Window System
- Many improvements since V2
- Device drivers run as user processes
- High degree of fault tolerance

The Book

- Andrew S Tanenbaum and Albert S Woodhull
- Operating Systems Design and Implementation, 3/E

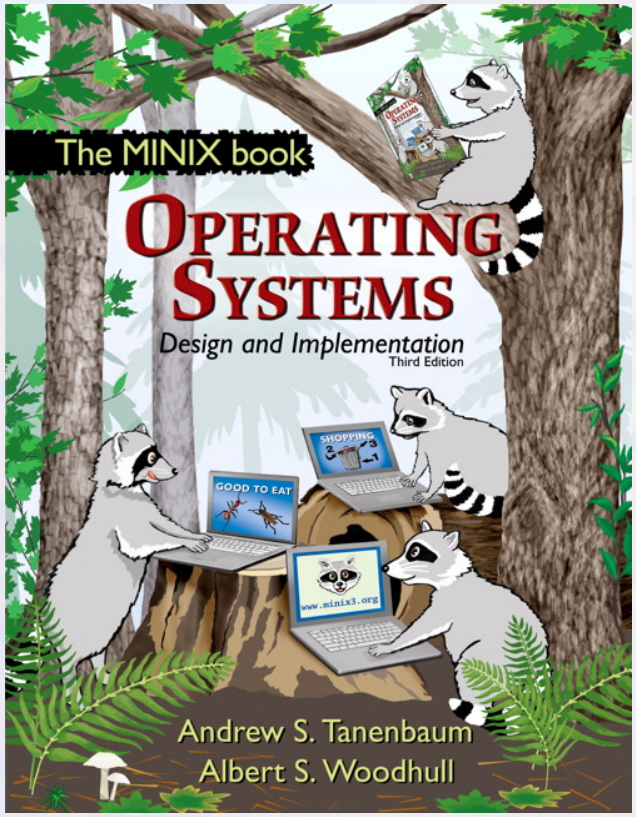

Minix System Calls (1)

System Calls (1)

- access - determine accessibility of file
- alarm - schedule signal after specified time
- brk, sbrk - change data segment size
- chdir, fchdir - change current working directory
- chmod - change mode of file
- chown - change owner and group of a file
- chroot - change root directory
- close - delete a descriptor
- creat - create a new file

Minix System Calls (2)

System Calls (2)

- close - delete a descriptor
- creat - create a new file
- dup, dup2 - duplicate a descriptor
- execve - execute a file
- exit,_exit - terminate a process
- fcntl - miscellaneous file descriptor control functions
- fork - create a new process
- getgid, getegid - get group identity
- getpid, getppid - get process identification

Minix System Calls (3)

System Calls (3)

- getpriority, setpriority - get and set scheduling priority
- gettimeofday - get date and time
- getuid, geteuid - get user identity
- intro, errno - introduction to system calls and error numbers
- ioctl - control device
- kill - send signal to a process
- link - make a hard link to a file
- lseek - move read/write pointer
- mkdir - make a directory file

Minix System Calls (4)

System Calls (4)

- mknod, mkfifo - make a special file
- mount, umount - mount or umount a file system
- open - open a file for reading or writing, or create a new file
- pause - stop until signal
- pipe - create an interprocess communication channel
- ptrace - process trace
- read - read input
- reboot - close down the system or reboot
- rename - change the name of a file

Minix System Calls (5)

System Calls (5)

- rmdir - remove a directory file
- select, FD_CLR, FD_ISSET, FD_SET, FD_ZERO - synchronous I/O multiplexing
- setsid, getpgrp - create process group, get process group id
- setuid, setgid - set user or group ID's
- sigaction, signal - manage signal state and handlers
- sigpending - report pending signals
- sigprocmask - manipulate the signal mask
- sigsuspend - suspend until signalled
- stat, lstat, fstat - get file status

Minix System Calls (6)

System Calls (6)

- svrctl - special server control functions
- sync, fsync - update dirty buffers and super-block
- time, stime - get/set date and time
- times - get process times
- umask - set file creation mode mask
- uname - get system info
- unlink - remove directory entry
- utime - set file times
- wait, waitpid - wait for process to terminate
- write - write output

Minix3 Software

- python-2.4.3 - python interpreter
- gcc-3.4.3 - GNU Compiler Collection v3.4.3
- gcc-4.1.1 - GCC 4.1.1, C and C++ compilers
- openssl-0.9.8a - library of security algorithms and protocols
- openssh-4.3p2 - openssh implementation of secure shell

Minix3 Software

and much more

`http://www.minix3.org/software/`

Minix3 News: SQLite

Wednesday 27 June 2007
SQLite is a small C library that implements a
self-contained, embeddable,
zero-configuration SQL database engine.

Some Papers about MINIX 3

- Construction of a Highly Dependable Operating System (in Proc. 6th European Dependable Comp. Conf., Oct 2006)
- Reorganizing UNIX for Reliability (in Proc. 11th ACSAC, Sept. 2006)
- MINIX 3: A Highly Reliable, Self-Repairing Operating System (in Oper. Sys. Rev., July 2006)
- Can We Make Operating Systems Reliable and Secure? (in IEEE Computer, May 2006, pp. 44-51)
- Modular System Programming in MINIX 3 (in USENIX ;login, April 2006, pp. 19-28)
- A Lightweight Method for Building Reliable Operating Systems Despite Unreliable Device Drivers (TR IR-CS-018)

VU Amsterdam Theses

- Design and Implementation of the MINIX Virtual File System
- A Port of the MINIX OS to the PowerPC Platform
- Building Performance Measurement Tools for the MINIX 3 Operating System
- Towards a True Microkernel Operating System

Other Projects

- Rewriting MINIX in Cyclone
- Dual core
- Performance benchmarking
- MINIX as multimedia server
- Measuring hotspot performance
- Port to $100 laptop Shared library support

Actual Research

`http://www.minix3.org/who_doing_what.html`

- New file system
- Kernel, FS, MM
- Reliability, OS architecture
- fault injection and testing
- USB driver
- Porting Samba
- Porting NFS

Porting MINIX 3 to Other Architectures

- Assembler for 64-bit AMD CPUs
- Chris Wade MIPS port
- Port to ARM

Minix3 on Xen

- MINIX on Xen
- `http://minixonxen.skynet.ie/cgi-bin/trac.cgi/wiki/Report`

Minix3 and VMware

Friday 27 October 2006
A ready-to-run Minix 3.1.2a for VMware Player
is available from the VMware page.

TFH-Berlin Diplomarbeit: Netfilter

- Diplomarbeit, Juli 2007, TFH Berlin 
- Brian Schüler
- Analysis and Porting of a network filtering architecture on Minix-3

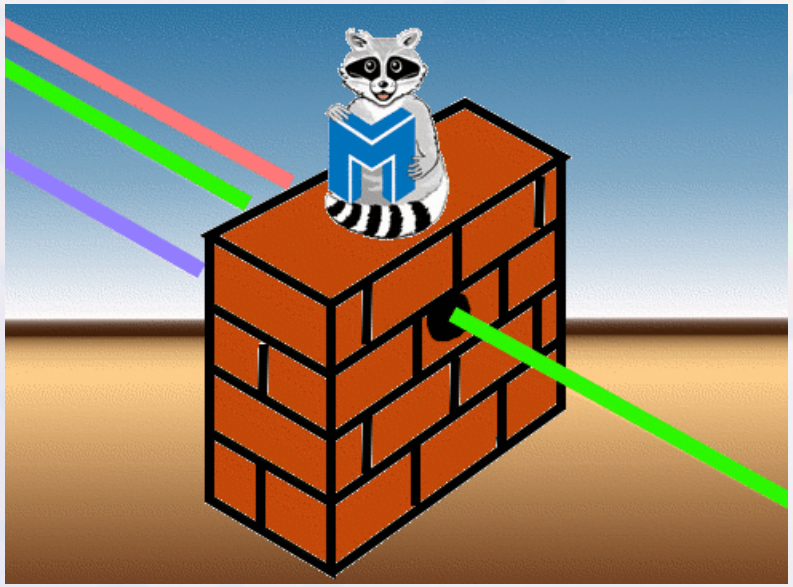

Minix3 Netfilter in USERMODE

	Linux Netfilter	Minix Netfilter
Crash Attack	System Crash	Restart Process
Executable Code	Owned System	Owned Usermode Process

MINIX 3 - Reliability

`http://www.minix3.org/reliability.html`

- Reduce kernel size
- Cage the bugs
- Limit drivers' memory access
- Restrict access to kernel functions
- Restrict access to I/O ports
- Restrict communication with OS components
- Reincarnate dead or sick drivers
- Survive bad pointers
- Tame infinite loops
- Limit damage from buffer overruns
- ...

Aktuelle Probleme im netfilter-Modul von Linux

-
  - 7.07.2007
  - Linux-Kernel-Update stopft ein Loch
  - `http://www.heise.de/newsticker/meldung/print/92369`
-
  - 22.03.2006
  - Buffer Overflow im netfilter-Modul von Linux
  - `http://www.heise.de/newsticker/meldung/print/71128`
-
  - 21.02.2005
  - Linux-Kernel-Patches beseitigen neue und alte Lücken
  - http://www.heise.de/newsticker/meldung/print/56625

Minix3 News: Wireless

Wednesday 1 Aug 2007

Driver for Orinoco wireless network cards

A driver for wireless PCI cards with the Prism chipset from Intersil by Michael Valkering and Stevens LeBlond has been committed.

USENIX ;login:

USENIX April 2006, Rick Farrow, "Musings"

`http://www.usenix.com/publications/login/2006-04/openpdfs/musings.pdf`

> *"While MINIX 3 is not going to replace your desktop today, it is already a good candidate for embedded systems where robustness, reliability, and a small memory footprint are crucial. Perhaps your cell phone will be running MINIX 3 some day."*

Disclamer

From: torvalds@klaava.Helsinki.FI (Linus Benedict Torvalds)
Subject: Re: LINUX is obsolete
Date: 29 Jan 92 23:14:26 GMT
Organization: University of Helsinki

''your job is being a professor and researcher:
That's one hell of a good excuse for some of
the brain-damages of minix.''

©opyleft

©opyleft

- Erstellt mit Freier Software
- © Rüdiger Weis, Berlin 2007
- unter der GNU Free Documentation License.