

CSC 405

Computer Security

Linux Security

Alexandros Kapravelos


akaprav@ncsu.edu

Unix / Linux

- Started in 1969 at AT&T / Bell Labs
- Split into a number of popular branches
 - BSD, System V (commercial, AT&T), Solaris, HP-UX, AIX
- Inspired a number of Unix-like systems
 - Linux, Minix, macOS
- Standardization attempts
 - POSIX, Single Unix Specification (SUS), Filesystem Hierarchy Standard (FHS), Linux Standard Base (LSB), ELF

OS Security

- Kernel vulnerability
 - usually leads to complete system compromise
 - attacks performed via system calls



Kernel vulnerabilities

#	CVE ID	CWE ID	# of Exploits	Vulnerability Type(s)	Publish Date	Update Date	Score	Gained Access Level	Access	Complexity	Authentication	Conf.	Integ.	Avail.
1	CVE-2017-12762	119		Overflow	2017-08-09	2017-08-25	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
In <code>drivers/isdn/l4/isdn_net.c</code> : A user-controlled buffer is copied into a local buffer of constant size using <code>strcpy</code> without a length check which can cause a buffer overflow. This affects the Linux kernel 4.9-stable tree, 4.12-stable tree, 3.18-stable tree, and 4.4-stable tree.														
2	CVE-2017-11176	416		DoS	2017-07-11	2017-08-07	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
The <code>mq_notify</code> function in the Linux kernel through 4.11.9 does not set the sock pointer to NULL upon entry into the retry logic. During a user-space close of a Netlink socket, it allows attackers to cause a denial of service (use-after-free) or possibly have unspecified other impact.														
3	CVE-2017-8890	415		DoS	2017-05-10	2017-05-24	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
The <code>inet_csk_clone_lock</code> function in <code>net/ipv4/inet_connection_sock.c</code> in the Linux kernel through 4.10.15 allows attackers to cause a denial of service (double free) or possibly have unspecified other impact by leveraging use of the accept system call.														
4	CVE-2017-7895	189			2017-04-28	2017-05-11	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
The NFSv2 and NFSv3 server implementations in the Linux kernel through 4.10.13 lack certain checks for the end of a buffer, which allows remote attackers to trigger pointer-arithmetic errors or possibly have unspecified other impact via crafted requests, related to <code>fs/nfsd/nfs3xdr.c</code> and <code>fs/nfsd/nfsxdr.c</code> .														
5	CVE-2017-0648	264		Exec Code	2017-06-14	2017-07-07	9.3	None	Remote	Medium	Not required	Complete	Complete	Complete
An elevation of privilege vulnerability in the kernel FIQ debugger could enable a local malicious application to execute arbitrary code within the context of the kernel. This issue is rated as High due to the possibility of a local permanent device compromise, which may require reflashing the operating system to repair the device. Product: Android. Versions: Kernel-3.10, Android ID: A-36101220.														
6	CVE-2017-0605	264		Exec Code	2017-05-12	2017-05-19	9.3	None	Remote	Medium	Not required	Complete	Complete	Complete
An elevation of privilege vulnerability in the kernel trace subsystem could enable a local malicious application to execute arbitrary code within the context of the kernel. This issue is rated as Critical due to the possibility of a local permanent device compromise, which may require reflashing the operating system to repair the device. Product: Android. Versions: Kernel-3.10, Kernel-3.18. Android ID: A-35399704. References: QC-CR#1048480.														
7	CVE-2017-0564	264		Exec Code	2017-04-07	2017-07-10	9.3	None	Remote	Medium	Not required	Complete	Complete	Complete
An elevation of privilege vulnerability in the kernel ION subsystem could enable a local malicious application to execute arbitrary code within the context of the kernel. This issue is rated as Critical due to the possibility of a local permanent device compromise, which may require reflashing the operating system to repair the device. Product: Android. Versions: Kernel-3.10, Kernel-3.18. Android ID: A-34276203.														
8	CVE-2017-0563	264		Exec Code	2017-04-07	2017-07-10	9.3	None	Remote	Medium	Not required	Complete	Complete	Complete
An elevation of privilege vulnerability in the HTC touchscreen driver could enable a local malicious application to execute arbitrary code within the context of the kernel. This issue is rated as Critical due to the possibility of a local permanent device compromise, which may require reflashing the operating system to repair the device. Product: Android. Versions: Kernel-3.10. Android ID: A-32089409.														
9	CVE-2017-0561	264		Exec Code	2017-04-07	2017-08-15	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
A remote code execution vulnerability in the Broadcom Wi-Fi firmware could enable a remote attacker to execute arbitrary code within the context of the Wi-Fi SoC. This issue is rated as Critical due to the possibility of remote code execution in the context of the Wi-Fi SoC. Product: Android. Versions: Kernel-3.10, Kernel-3.18. Android ID: A-34199105. References: B-RB#110814.														
10	CVE-2017-0528	264		Exec Code Bypass	2017-03-07	2017-07-17	9.3	None	Remote	Medium	Not required	Complete	Complete	Complete
An elevation of privilege vulnerability in the kernel security subsystem could enable a local malicious application to execute code in the context of a privileged process. This issue is rated as High because it is a general bypass for a kernel level defense in depth or exploit mitigation technology. Product: Android. Versions: Kernel-3.18. Android ID: A-33351919.														

Kernel vulnerabilities

#	CVE ID	CWE ID	# of Exploits	Vulnerability Type(s)	Publish Date	Update Date	Score	Gained Access Level	Access	Complexity	Authentication	Conf.	Integ.	Avail.
1	CVE-2018-20961	415		DoS	2019-08-07	2019-08-27	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
In the Linux kernel before 4.16.4, a double free vulnerability in the f_midi_set_alt function of drivers/usb/gadget/function/f_midi.c in the f_midi driver may allow attackers to cause a denial of service or possibly have unspecified other impact.														
2	CVE-2019-10125	94			2019-03-27	2019-06-14	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
An issue was discovered in aio_poll() in fs/aio.c in the Linux kernel through 5.0.4. A file may be released by aio_poll_wake() if an expected event is triggered immediately (e.g., by the close of a pair of pipes) after the return of vfs_poll(), and this will cause a use-after-free.														
3	CVE-2019-11583	399		DoS Mem. Corr.	2019-05-02	2019-06-14	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
udp_gro_receive_segment in net/ipv4/udp_offload.c in the Linux kernel 5.x before 5.0.13 allows remote attackers to cause a denial of service (slab-out-of-bounds memory corruption) or possibly have unspecified other impact via UDP packets with a 0 payload, because of mishandling of padded packets, aka the "GRO packet of death" issue.														
4	CVE-2019-11811	416			2019-05-07	2019-05-31	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
An issue was discovered in the Linux kernel before 5.0.4. There is a use-after-free upon attempted read access to /proc/ioports after the ipmi_si module is removed, related to drivers/char/ipmi/ipmi_si_intf.c, drivers/char/ipmi/ipmi_si_mem_io.c, and drivers/char/ipmi/ipmi_si_port_io.c.														
5	CVE-2019-15292	416			2019-08-21	2019-09-02	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
An issue was discovered in the Linux kernel before 5.0.9. There is a use-after-free in atalk_proc_exit, related to net/appletalk/atalc_proc.c, net/appletalk/ddp.c, and net/appletalk/sysctl_net_atalk.c.														
6	CVE-2019-15504	415			2019-08-23	2019-09-04	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
drivers/net/wireless/rsi/rsi_91x_usb.c in the Linux kernel through 5.2.9 has a Double Free via crafted USB device traffic (which may be remote via usbip or usbredir).														
7	CVE-2019-15505	125			2019-08-23	2019-09-04	10.0	None	Remote	Low	Not required	Complete	Complete	Complete
drivers/media/usb/dvb-usb/technisat-usb2.c in the Linux kernel through 5.2.9 has an out-of-bounds read via crafted USB device traffic (which may be remote via usbip or usbredir).														
8	CVE-2019-15926	125			2019-09-04	2019-09-14	9.4	None	Remote	Low	Not required	Complete	None	Complete
An issue was discovered in the Linux kernel before 5.2.3. Out of bounds access exists in the functions ath6kl_wmi_pstream_timeout_event_rx and ath6kl_wmi_cac_event_rx in the file drivers/net/wireless/ath/ath6kl/wmi.c.														
9	CVE-2018-20836	416			2019-05-07	2019-05-08	9.3	None	Remote	Medium	Not required	Complete	Complete	Complete
An issue was discovered in the Linux kernel before 4.20. There is a race condition in smp_task_timeout() and smp_task_done() in drivers/scsi/libsas/sas_expander.c, leading to a use-after-free.														
10	CVE-2019-18115	362			2019-05-08	2019-06-07	9.3	None	Remote	Medium	Not required	Complete	Complete	Complete
An issue was discovered in rds_tcp_kill_sock in net/rds/tcp.c in the Linux kernel before 5.0.8. There is a race condition leading to a use-after-free, related to net namespace cleanup.														

Kernel security research is active

- Kernel Security track on USENIX Security 2022
 - Playing for K(H)eaps: Understanding and Improving Linux Kernel Exploit Reliability
 - In-Kernel Control-Flow Integrity on Commodity OSes using ARM Pointer Authentication
 - Midas: Systematic Kernel TOCTTOU Protection
 - LinkRID: Vetting Imbalance Reference Counting in Linux kernel with Symbolic Execution

Kernel security is rapidly changing

Rust will go into Linux 6.1

random example

Unix

- Code running in user mode is always linked to a certain identity
 - security checks and access control decisions are based on user identity
- Unix is user-centric
 - no roles
- User
 - identified by username (UID), group name (GID)
 - typically authenticated by password (stored encrypted)
- User root
 - superuser, system administrator
 - special privileges (access resources, modify OS)
 - cannot decrypt user passwords

Process Management

- Process
 - implements user-activity
 - entity that executes a given piece of code
 - has its own execution stack, memory pages, and file descriptors table
 - separated from other processes using the virtual memory abstraction
- Thread
 - separate stack and program counter
 - share memory pages and file descriptor table

Process Management

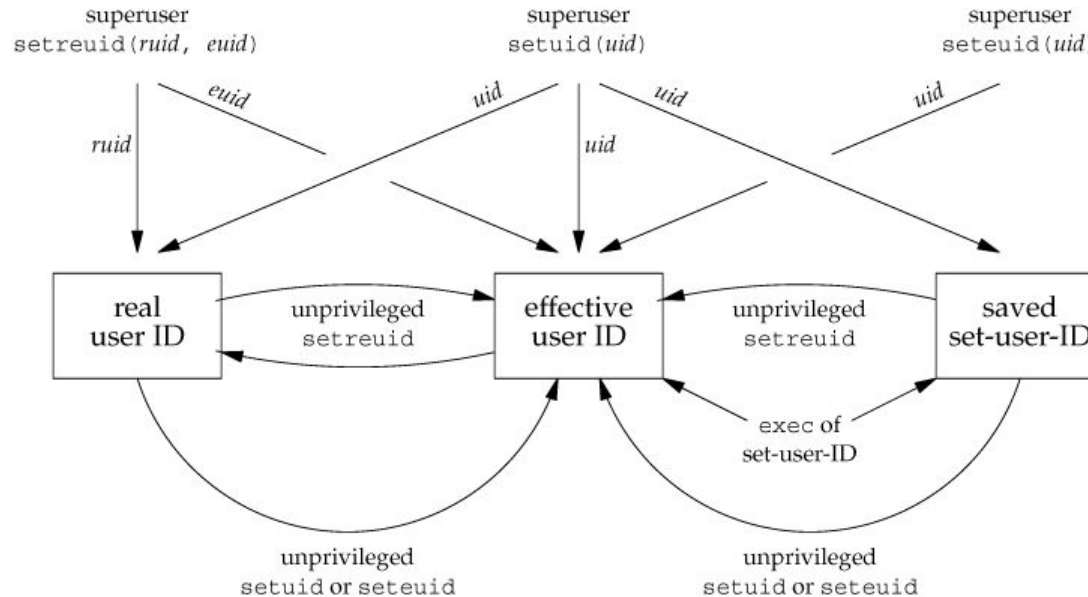
- Process Attributes
 - process ID (PID)
 - uniquely identified process
 - (real) user ID (UID)
 - ID of owner of process
 - effective user ID (EUID)
 - ID used for permission checks (e.g., to access resources)
 - saved user ID (SUID)
 - to temporarily drop and restore privileges
 - lots of management information
 - scheduling
 - memory management, resource management

Process Management

- Switching between IDs
 - uid-setting system calls

```
int setuid(uid_t uid)
int seteuid(uid_t uid)
int setresuid(uid_t ruid, uid_t euid, uid_t suid)
```
- Can be tricky
 - POSIX 1003.1:
If the process has appropriate privileges, the setuid(newuid) function sets the real user ID, effective user ID, and the [saved user ID] to newuid.
 - what are appropriate privileges?
Solaris: EUID = 0; FreeBSD: newuid = EUID;
Linux: SETUID capability

Summary of all the functions that set the various user IDs



Process Management

Bug in sendmail 8.10.1:

- call to `setuid(getuid())` to clear privileges (effective UID is root)
- on Linux, attacker could clear SETUID capability
- call clears EUID, but SUID remains root

Further reading

Setuid Demystified

Hao Chen, David Wagner, and Drew Dean

11th USENIX Security Symposium, 2002

User Authentication

- How does a process get a user ID?
- Authentication
- Passwords
 - user passwords are used as keys for crypt() function
 - uses SHA-512
 - 8-byte “salt”
 - chosen from date, not secret
 - prevent same passwords to map onto same string
 - make dictionary attacks more difficult
- Password cracking
 - dictionary attacks, rainbow tables
 - Crack, JohnTheRipper

User Authentication

- Shadow passwords
 - password file is needed by many applications to map user ID to user names
 - encrypted passwords are not
- /etc/shadow
 - holds encrypted passwords
 - account information
 - last change date
 - expiration (warning, disabled)
 - minimum change frequency
 - readable only by superuser and privileged programs
 - SHA-512 hashed passwords (default on Ubuntu) to slow down guessing

User Authentication

- Shadow passwords
 - a number of other encryption / hashing algorithms were proposed
 - blowfish, SHA-1, ...
- Other authentication means possible
 - Linux PAM (pluggable authentication modules)
 - Kerberos
 - Active directory (Windows)

Group Model

- Users belong to one or more groups
 - primary group (stored in /etc/passwd)
 - additional groups (stored in /etc/group)
 - possibility to set group password
 - and become group member with newgrp

- /etc/group

```
groupname : password : group id : additional users  
root:x:0:root  
bin:x:1:root,bin,daemon  
users:x:100:akaprav
```

- Special group wheel/sudo
 - protect root account by limiting user accounts that can perform su

File System

- File tree
 - primary repository of information
 - hierarchical set of directories
 - directories contain file system objects (FSO)
 - root is denoted “/”
- File system object
 - files, directories, symbolic links, sockets, device files
 - referenced by inode (index node)

File System

- Access Control
 - permission bits
 - `chmod`, `chown`, `chgrp`, `umask`
 - file listing:

 - **rwX** **rwX** **rwX**
(file type) (user) (group) (other)

Type	r	w	x	s	t
File	read access	write access	execute	suid / sgid inherit id	sticky bit
Directory	list files	insert and remove files	stat / execute files, chdir	new files have dir-gid	files/dirs only delete-able by owner

Sticky bit

- It has no effect on files (on Linux)
- When used on a directory, all the files in that directory will be modifiable only by their owners
- What's a very common directory with sticky bit?

```
$ ls -ld /tmp
```

```
drwxrwxrwt 26 root root 69632 Sep  7 15:24 /tmp
```

```
$ ls -l test
```

```
-rw-rw-r-- 1 kapravel kapravel 0 Sep  7 15:29 test
```

```
$ chmod +t test; ls -l test
```

```
-rw-rw-r-T 1 kapravel kapravel 0 Sep  7 15:29 test
```

SUID Programs

- Each process has real and effective user / group ID
 - usually identical
 - real IDs
 - determined by current user
 - authentication (login, su)
 - effective IDs
 - determine the “rights” of a process
 - system calls (e.g., `setuid()`)
 - suid / sgid bits
 - to start process with effective ID different from real ID
 - attractive target for attacker
- Never use SUID shell scripts (multiplying problems)

File System

- Shared resource
 - susceptible to race condition problems
- Time-of-Check, Time-of-Use (**TOCTOU**)
 - common race condition problem
 - problem:
 - Time-Of-Check (t_1): validity of assumption A on entity E is checked
 - Time-Of-Use (t_2): assuming A is still valid, E is used
 - Time-Of-Attack (t_3): assumption A is invalidated

$$t_1 < t_3 < t_2$$

TOCTOU

- Steps to access a resource
 - obtain reference to resource
 - query resource to obtain characteristics
 - analyze query results
 - if resource is fit, access it
- Often occurs in Unix file system accesses
 - check permissions for a certain file name (e.g., using `access(2)`)
 - open the file, using the file name (e.g., using `fopen(3)`)
 - four levels of indirection (symbolic link - hard link - inode - file descriptor)
- Windows uses file handles and includes checks in API open call

Overview

```
/* access returns 0 on success */  
if(!access(file, W_OK)) {  
    f = fopen(file, "wb+");  
    write_to_file(f);  
} else {  
    fprintf(stderr, "Permission denied \  
                    when trying to open %s.\n", file);  
}
```

- Attack

```
$ touch dummy; ln -s dummy pointer  
$ rm pointer; ln -s /etc/passwd pointer
```

Examples

- TOCTOU Examples

- Setuid Scripts

- `exec()` system call invokes `setuid()` call prior to executing program
 - program is a script, so command interpreter is loaded first
 - program interpreted (with root privileges) is invoked on script name
 - attacker can replace script content between step 2 and 3

Examples

- TOCTOU Examples
 - Directory operations
 - `rm` can remove directory trees, traverses directories depth-first
 - issues `chdir("..")` to go one level up after removing a directory branch
 - by relocating subdirectory to another directory, arbitrary files can be deleted
 - Temporary files
 - commonly opened in `/tmp` or `/var/tmp`
 - often guessable file names

Temporary Files

- “Secure” procedure for creating temporary files
 - pick a prefix for your filename
 - generate at least 64 bits of high-quality randomness
 - base64 encode the random bits
 - concatenate the prefix with the encoded random data
 - set umask appropriately (0066 is usually good)
 - use `fopen(3)` to create the file, opening it in the proper mode
 - delete the file immediately using `unlink(2)`
 - perform reads, writes, and seeks on the file as necessary
 - finally, close the file

Temporary Files

- Library functions to create temporary files can be insecure
 - `mktemp(3)` is not secure, use `mkstemp(3)` instead
 - old versions of `mkstemp(3)` did not set `umask` correctly
- Temp Cleaners
 - programs that clean “old” temporary files from temp directories
 - first `lstat(2)` file, then use `unlink(2)` to remove files
 - vulnerable to race condition when attacker replaces file between `lstat(2)` and `unlink(2)`
 - arbitrary files can be removed
 - delay program long enough until temp cleaner removes active file

Prevention

- Immutable bindings
 - operate on file descriptors
 - do not check access by yourself (i.e., no use of `access(2)`)
drop privileges instead and let the file system do the job
- Use the `O_CREAT` | `O_EXCL` flags to create a new file with `open(2)`
 - be prepared to have the `open` call fail

Prevention

Series of papers on the access system call

Fixing races for fun and profit: how to use access(2)

D. Dean and A. Hu

Usenix Security Symposium, 2004

Fixing races for fun and profit: howto abuse atime

N. Borisov, R. Johnson, N. Sastry, and D. Wagner

Usenix Security Symposium, 2005

Portably Solving File TOCTTOU Races with Hardness Amplification

D. Tsafir, T. Hertz, D. Wagner, and D. Da Silva

Usenix Conference on File and Storage Technologies (FAST), 2008

Locking

- Ensures exclusive access to a certain resource
- Used to circumvent accidental race conditions
 - advisory locking (processes need to cooperate)
 - not mandatory, therefore not secure
- Often, files are used for locking
 - portable (files can be created nearly everywhere)
 - “stuck” locks can be easily removed
- Simple method
 - create file using the O_EXCL flag

Shell

- Shell
 - one of the core Unix application
 - both a command language and programming language
 - provides an interface to the Unix operating system
 - rich features such as control-flow primitives, parameter passing, variables, and string substitution
 - communication between shell and spawned programs via redirection and pipes
 - different flavors
 - bash and sh, tcsh and csh, ksh, zsh

Shell Attacks

- Environment Variables
 - \$HOME and \$PATH can modify behavior of programs that operate with relative path names
 - \$IFS – internal field separator
 - used to parse tokens
 - usually set to [\t\n] but can be changed to “/”
 - “/bin/ls” is parsed as “bin ls” calling bin locally
 - IFS now only used to split expanded variables
 - preserve attack (/usr/lib/preserve is SUID)
 - called “/bin/mail” when vi crashes to preserve file
 - change IFS, create bin as link to /bin/sh, kill vi

Shell Attacks

- Control and escape characters
 - can be injected into command string
 - modify or extend shell behavior
 - user input used for shell commands has to be rigorously sanitized
 - easy to make mistakes
 - classic examples are `;` and `&`
- Applications that are invoked via shell can be targets as well
 - increased vulnerability surface
- Restricted shell
 - invoked with `-r` or `rbash`
 - more controlled environment

Shell Attacks

- `system(char *cmd)`
 - function called by programs to execute other commands
 - invokes shell
 - executes string argument by calling `/bin/sh -c string`
 - makes binary program vulnerable to shell attacks
 - especially when user input is utilized
- `popen(char *cmd, char *type)`
 - forks a process, opens a pipe and invokes shell for cmd

File Descriptor Attacks

- SUID program opens file
- forks external process
 - sometimes under user control
- on-execute flag
 - if close-on-exec flag is not set, then new process inherits file descriptor
 - malicious attacker might exploit such weakness
- Linux Perl 5.6.0
 - `getpwuid()` leaves `/etc/shadow` opened (June 2002)
 - problem for Apache with `mod_perl`

Resource Limits

- File system limits
 - quotas
 - restrict number of storage blocks and number of inodes
 - hard limit
 - can never be exceeded (operation fails)
 - soft limit
 - can be exceeded temporarily
 - can be defined per mount-point
 - defend against resource exhaustion (denial of service)
- Process resource limits
 - number of child processes, open file descriptors

Signals

- Signal
 - simple form of interrupt
 - asynchronous notification
 - can happen anywhere for process in user space
 - used to deliver segmentation faults, reload commands, ...
 - kill command
- Signal handling
 - process can install signal handlers
 - when no handler is present, default behavior is used
 - ignore or kill process
 - possible to catch all signals except SIGKILL (-9)

Signals

- Security issues
 - code has to be re-entrant
 - atomic modifications
 - no global data structures
 - race conditions
 - unsafe library calls, system calls
 - examples
 - wu-ftpd 2001, sendmail 2001 + 2006, stunnel 2003, ssh 2006
- Secure signals
 - write handler as simple as possible
 - block signals in handler

Shared Libraries

- Library
 - collection of object files
 - included into (linked) program as needed
 - code reuse
- Shared library
 - multiple processes share a single library copy
 - save disk space (program size is reduced)
 - save memory space (only a single copy in memory)
 - used by virtually all Unix applications (at least libc.so)
 - check binaries with ldd

Shared Libraries










- Static shared library
 - address binding at link-time
 - not very flexible when library changes
 - code is fast
- Dynamic shared library
 - address binding at load-time
 - uses procedure linkage table (PLT) and global offset table (GOT)
 - code is slower (indirection)
 - loading is slow (binding has to be done at run-time)
 - classic .so or .dll libraries
- PLT and GOT entries are very popular attack targets
 - buffer overflows

Shared Libraries

- Management
 - stored in special directories (listed in `/etc/ld.so.conf`)
 - manage cache with `ldconfig`
- Preload
 - override (substitute) with other version
 - use `/etc/ld.so.preload`
 - can also use environment variables for override
 - possible security hazard
 - now disabled for SUID programs (old Solaris vulnerability)

Advanced Security Features

- Address space protection
 - address space layout randomization (ASLR)
 - non-executable stack (based on NX bit or PAX patches)
- Mandatory access control extensions
 - SELinux/AppArmor
 - role-based access control extensions
 - capability support
- Miscellaneous improvements
 - hardened chroot jails
 - better auditing



3 Receiver Antennas
amera for Annotation

