

# **Linux Kernel == Security Nightmare**

**Marcel Holtmann**

**Red Hat Security Response Team**

*EUSecWest 2007 Conference: London, England*

# It's an expression

- True
  - Please wake up
  - We are in big trouble and it is time to do something
- False
  - Don't fall asleep
  - We have to make sure that it stays this way

# Agenda

- The Linux kernel
- Some words about security response
- Difference between upstream and distributions/vendors
- Deep look at some vulnerabilities in the Linux kernel from the last 6 month
- Technologies to better secure your systems

# Some statistics

- Linux 2.6.19 kernel
  - 17.133 files
  - 6.999.908 lines of code (3.488.831 for drivers)
- Firefox 1.5.0.7
  - 6.386 files
  - 1.837.531 lines of code

# Security response

- Handle security issues in time
- Research possible impacts
- Determine affected versions
- Assign CVE name
- Communicate with other vendors
- Handle embargoes
- Release updates

# Information sources

- [vendor-sec@lst.de](mailto:vendor-sec@lst.de)
  - Closed group of security experts from vendors
  - Mainly Linux (Unix) based and invitation only
- [security@kernel.org](mailto:security@kernel.org)
  - Small group of around 6 people from the Linux kernel community
- Full disclosure and Bugtraq
  - Public mailing lists
- Reports from researchers

# Severity level

- Red Hat uses the same classification scheme as Microsoft does

<http://www.redhat.com/security/updates/classification/>



# Severities in detail

- **Critical** *“A vulnerability whose exploitation could allow the propagation of an Internet worm without user action.”*
- **Important** *“easily compromise the Confidentiality, Integrity or Availability of resources”*
- **Moderate** *“harder or more unlikely to be exploitable”*
- **Low** *“unlikely circumstances ... or where a successful exploit would lead to minimal consequences”*

# Affected kernels

- Upstream kernel
  - Mainline 2.4 and 2.6 kernels
  - Stable branches
  - The -mm kernel
- Distribution kernel
  - Branched from upstream kernel
  - Backported patches and features
  - The Red Hat Enterprise Linux 2.1 kernel is based on 2.4.9 (August, 16<sup>th</sup> 2001)

# CVE names

- Common Vulnerabilities and Exposures

*“A list of standardized names for vulnerabilities and other information security exposures — CVE aims to standardize the names for all publicly known vulnerabilities and security exposures.”*

<http://cve.mitre.org/>

- Example: CVE-2006-2451

# NVD database

- Database for vendors based on CVE names
- Possibility to give vendor statements

<http://nvd.nist.gov/>

- Red Hat uses it for official statement of security flaws that doesn't affect their products

# Embargoes

- Different opinion from different people
  - Everybody can have his own opinion
- Sensible use of embargoes
  - Needed to keep the days of risk minimal
  - Balance between customers and open source
  - General embargo time is 1-2 weeks
  - Release only from Tuesday to Thursday
  - Communicate through vendor-sec

# Release policy

- **Critical vulnerabilities**
  - Will be pushed immediately an embargo is lifted, or when passed QE
  - Will be pushed at any time or day
- **Important vulnerabilities**
  - May be held until reasonable time or day
- **Moderate or low vulnerabilities**
  - May be held until other issues come up in the same package, or the next update release

# Kernel update cycle

- Upstream normally releases a new kernel version every 3 month
- Stable kernels are released at will, but mostly for security reasons
- Security updates for a distribution kernel in general only once a month

# Categorize vulnerabilities

- Privilege escalation
  - Gain root access
- Denial of service (local and remote)
  - In form of panics, crashes etc.
- Information leaks
  - Access memory areas with sensible data

# Problematic areas

- The netfilter code
  - Needed for firewalling etc.
- New network protocols
  - For example IPv6 or SCTP
- Not widely used architectures
  - Machines with PowerPC or UltraSparc CPUs
- Filesystems to some degree

# Virtualization efforts

- Hypervisor is the new magic word
- Generic paravirt\_ops interface
  - Xen, VMware, KVM, kQEMU
  - OpenVZ
- Malicious guests are a problem
- Host system should be flawless
- The real nightmare

# CVE-2006-1864

- Breaking chroot on SMB share
- Affects smbfs and cifs
- The 2.4 and 2.6 kernels were vulnerable
  
- In case of cifs the backport was ugly
- Use of chroot with SMB is unlikely

# CVE-2006-2274

- Remote denial of service attack against the SCTP stack
- Caused an infinite recursion and will stall the system
- Only one of the possible SCTP issues

# CVE-2006-0457

- Denial of service or information leak in keyring handling
- Non privileged users were able to crash the kernel
- Possible to retrieve sensitive information about encrypted filesystems etc.

# CVE-2006-4813

- Information leak
- The function `__block_prepare_write()` didn't clear its used memory
- Possible to read root-only files
- Leaking serious amount of data

# CVE-2006-3468

- Bogus NFS request caused denial of service
- The ext3 filesystem shuts down and mounts itself read-only
- Incorrect handling of error cases made ext3 vulnerable

# CVE-2006-2451

- Privilege escalation through `prctl()`
- Basically a design flaw
- Embargoed for 2 weeks
- Used to break into Debian and Sourceforge servers
- Red Hat provided updated kernel on the date of publication

# CVE-2006-3626

- Privilege escalation through /proc
- Race condition and design flaw
- 0-day exploit on a Friday evening
- Fixed upstream within 6 hours
- SELinux default policy prevented the exploitation on RHEL4
- The 2.4 kernel series was not affected

# CVE-2006-5753

- Wrong user vs. kernel space pointer
- Easily overwriting of user memory
  
- Requires a bad inode to trigger the flaw

# CVE-2007-0004

- ACL handling for NFS shares
- Possible to open and read files from other users (including root)
- Wrong caching of ACL information
  
- Linux 2.4 only issue
- Code has been backported

# CVE-2007-0006

- Key serial number collision flaw
  - Overwriting sensible kernel memory
  - Incremented vs. random serial numbers
  - 2 Billions keys before crashing when using incremented serial numbers
- 
- Random is not really random!

# **CVE-2006-3635**

- This issue is still embargoed

# Issue overview

- Most issues are local denial of services
  - Minor important if no local or untrusted users exists on the system
- Remote denial of services happens
  - Serious if no firewall or other protection is in place to secure the system
- Privilege escalation / information leaks
  - Posing a real thread for systems with local or untrusted users

# Upstream effort

- Creation of the -stable kernel series
- Supports the last two kernel releases
- Fast response to security issues
- Assigning CVE names for all issues
  
- Maintained by Greg Kroah-Hartman from SuSE/Novell and Chris Wright from Red Hat

# Kernel releases

- Linux 2.6.16
    - Linux 2.6.16.42
  - Linux 2.6.17
    - Linux 2.6.17.14
  - Linux 2.6.18
    - Linux 2.6.18.8
  - Linux 2.6.20
    - Should deprecate the 2.6.18 stable series
- 2006-Mar-20  
2007-Feb-26
- 2006-Jun-18  
2006-Oct-13
- 2006-Sep-20  
2007-Feb-23
- 2007-Feb-04

# Conclusion

- Kernel security is taken serious
  - Sensible embargoes
  - Excellent 0-day response time
  - SELinux and Exec-shield helps
  - Keep userspace under control
- 
- And yes, it is not perfect ... but we are trying hard to make it better every day

# Thanks

- Have a good night sleep and dream something nice ...

