

Linux Kernel == Security Nightmare

Marcel Holtmann

Red Hat Security Response Team

PacSec 2006 Conference: Tokyo, Japan

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

- Some words about security response for the Linux kernel
- Insights to our processes
- 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

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**
 - Closed group of security experts from vendors
 - Mainly Linux (Unix) based
 - Invitation only
- **security@kernel.org**
 - Small group of around 6 people from the Linux kernel community
- **Full disclosure and Bugtraq**
 - Public mailing lists

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
- 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, 16th 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

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 months
- 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

CVE-2006-1864

- Breaking chroot on SMB share
- Affects smbfs and cifs
- The 2.4 and 2.6 kernels are 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
- Causes 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 user could crash the kernel
- Possible to retrieve sensitive information about encrypted filesystems etc.

CVE-2006-4813

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

CVE-2006-3468

- Bogus NFS request causes denial of service
- The ext3 filesystem shuts down and mounts itself read-only
- Incorrect handling of error cases makes 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
- 0day 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-2006-3635

- 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

Red Hat innovations

- Reducing attack vectors
 - 2001 Firewall on by default
 - 2004 NX and software NX by default
 - 2004 Randomization
 - 2005 Heap overflow checks
 - 2005 SELinux on by default
 - 2006 GLibc and GCC checks
- Keep the user space under control

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

Conclusion

- Kernel security is taken serious
 - Sensible embargoes
 - Excellent 0day response time
 - SELinux and Exec-shield helps
-
- 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 ...

