# Practical Grsecurity in a nutshell

# whoami

- @k0ck4 – kocka was already taken on twitter
- Grsecurity and Alpine Linux evangelist
- sysadmin / system engineer
- Itsec wannabe

# What is Grsecurity 1.

- “ Grsecurity® is an extensive security enhancement to the Linux kernel that defends against a wide range of security threats through intelligent access control, memory corruption-based exploit prevention, and a host of other system hardening that generally require no configuration. It has been actively developed and maintained for the past 14 years. Commercial support for grsecurity is available through Open Source Security, Inc. “

# What is Grsecurity 2.

- PaX (memory-corruption based defenses, leak prevention)
  - MPROTECT
  - UDEREF
  - PAGEEXEC / SEGMEXEC
  - ASLR
  - gcc plugins (eg: size overflow, stackleak, constify, latent_entropy)
  - randmmap
  - randkstack, randustack
  - etc..
- RBAC
- Other hardening, kernel audit, continuous backport of fixes with security impact from patches sent to upstream

# PaX influenced projects

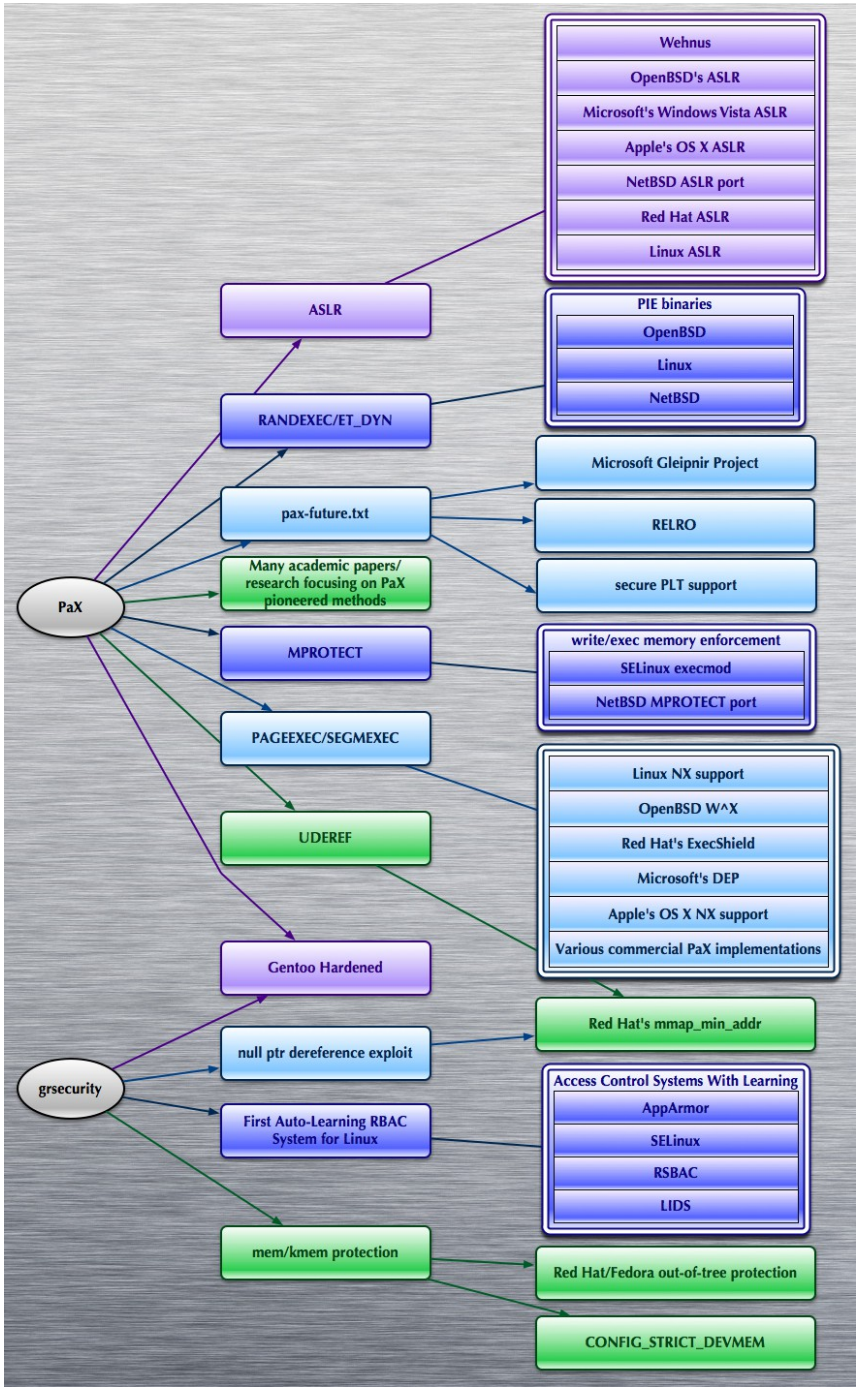

# ASLR

- Upstream don't merge PaX ASLR, implement it on their own way:
  - https://twitter.com/grsecurity/status/542316973963878400
  - http://marc.info/?l=linux-kernel&m=141911002822659&w=2
  - http://seclists.org/oss-sec/2014/q4/986
  - https://lkml.org/lkml/2015/2/14/61
  - http://cybersecurity.upv.es/attacks/offset2lib/offset2lib.html

    “Pax solution also increases the number of entropy of each zone, even it is able to randomise non-PIE applications. As far as we know it is the most advanced ASLR implementation. Unfortunately, some people think that it is a too complex patch with, may be, too many features (some advanced features may break backward compatibility on some applications).”

# RBAC

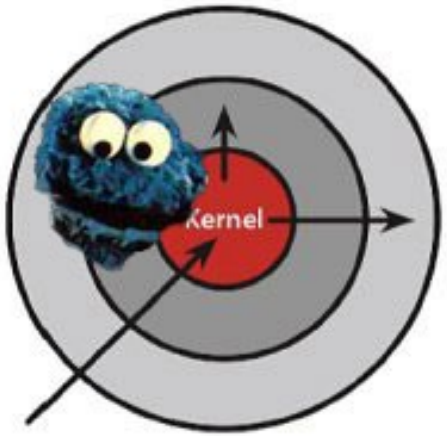


**Discretionary Access Control**

Once a security exploit gains access to priveleged system component, the entire system is compromised.

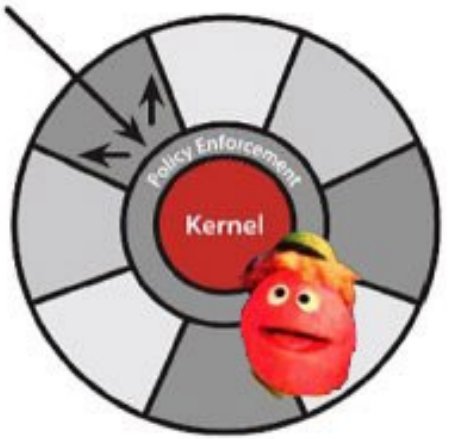


**Mandatory Access Control**

Kernel policy defines application rights, firewalling applications from compromising the entire system.

**Blackhats with kernel exploits**

Basement dwelling 12-year olds armed with kernel exploit released past Tuesday. A SELinux disabling payload in the exploit turns your entire MAC policy into laughing stock. You spend the rest of the weekend removing SSH backdoors.

Red Hat and Security-Enhanced Linux (SELinux): It's really about the neat diagrams.

# Linus's attitude 1.

- Security bugs are just bugs, don't want to participate in the security circus (source: http://article.gmane.org/gmane.linux.kernel/706950 )

  "...one reason I refuse to bother with the whole security circus is that I think it glorifies -- and thus encourages -- the wrong behavior. It makes 'heroes' out of security people, as if the people who don't just fix normal bugs aren't as important. In fact, all the boring normal bugs are way more important, just because there's a lot more of them.
  ...I think the OpenBSD crowd is a bunch of masturbating monkeys, in that they make such a big deal about concentrating on security to the point where they pretty much admit that nothing else matters to them."

# Linus's attitude 2.

- Silently fix bugs, not mention security impact in commit logs (eg: http://arstechnica.com/security/2013/05/critical-linux-vulnerability-imperils-users-even-after-silent-fix/ )
- It would be the linux distributions maintainers task to verify every single patch to the upstream?
- Fortunately Spender takes care:
- 2015.06.03: “Committed upstream today 4 mo later,grsec had it Feb 2nd” https://twitter.com/grsecurity/status/606226550183325697

# Wrestling with upstream developers

- “Just use a supported kernel.”
- “Stop reopening the bug.  If you have to redefine what your own OS does then maintain your own version of everything else as we”
  - Ulrich Drepper

    Source: https://sourceware.org/bugzilla/show_bug.cgi?id=12492

# DEMO

# DEMO 1.

- Jul 10 00:36:41 shannon kernel: [ 6542.009341] grsec: denied RWX mprotect of <anonymous mapping> by /usr/lib/libreoffice/program/soffice.bin[soffice.bin:7151] uid/euid:

  1000/1000 gid/egid:1001/1001, parent /usr/lib/libreoffice/program/oosplash[oosplash:7150] uid/euid:1000/1000 gid/egid:1001/1001

- root@shannon:/home/kocka# setfattr -n user.pax.flags -v "m" /usr/lib/libreoffice/program/soffice.bin
- root@shannon:/home/kocka# getfattr -n user.pax.flags /usr/lib/libreoffice/program/soffice.bin

  getfattr: Removing leading '/' from absolute path names

  # file: usr/lib/libreoffice/program/soffice.bin

  user.pax.flags="m"
- root@shannon:/home/kocka# setfattr -x user.pax.flags /usr/lib/libreoffice/program/soffice.bin
- root@shannon:/home/kocka# getfattr -n user.pax.flags /usr/lib/libreoffice/program/soffice.bin

  /usr/lib/libreoffice/program/soffice.bin: user.pax.flags: No such attribute

# DEMO 2.

- kocka@shannon:~/campp$ ./script.sh

  bash: ./script.sh: /bin/sh: bad interpreter: Permission denied
- kernel: [15463.642991] grsec: denied untrusted exec (due to not being in trusted group and file in non-root-owned directory) of /home/kocka/campp/script.sh by /home/kocka/campp/script.sh[bash:22712] uid/euid:1000/1000 gid/egid:1001/1001, parent /bin/bash[bash:14743] uid/euid:1000/1000 gid/egid:1001/1001

# Related projects (a few)

- Copperhead OS / https://copperhead.co/2015/06/11/android-pax
- HardenedBSD
- Alpine Linux (contains grsec kernel by default)
- hardened gentoo
- Corsac's .deb repository
- mempo - reproducible builds vs. randomization

# Summary

- Grsecurity is a way too cool project, it's a sin not using it. Keep up the good work!
- Sooner or later you will find a bug
  - Don't disable grsec and walk away, report it and help debugging which component is buggy (crashed software, vanilla kernel, grsec)
  - Eg:
    - https://bugs.alpinelinux.org/issues/3061
    - https://forums.grsecurity.net/viewtopic.php?f=3&t=3977
    - https://bugs.archlinux.org/task/40627

```
commit b4a3ab65850c171ca72716ad05a39d16158e45e4
Author: Brad Spengler <spender@grsecurity.net>
Date:   Sat Jun 21 23:17:23 2014 -0400

    Fix GRKERNSEC_KSTACKOVERFLOW incompatibility with virtio_net and other more
    rare drivers.  Unfortunately to resolve the problem we had to choose between
    invasive changes to dozens of call-sites and continued future maintenance work,
    or rearchitecting the feature to be able to handle the uses seamlessly.  With
    some tips from pipacs, I chose the latter.

    Various drivers including virtio_net use scatterlists derived from stack-based
    buffers (e.g. as an argument to sg_set_buf/sg_init_one).  The scatterlist API
    requires that these buffers be in the kernel image or in kmalloc'd buffers,
    which caused a problem when vmalloc'd stacks were used due to
    GRKERNSEC_KSTACKOVERFLOW.  What we do now is keep the original lowmem kstack
    allocation and then perform a THREAD_SIZE-aligned vmapped alias of the lowmem
    kstack's physical pages.  We also restore kernel stack accounting by using
    this method.  The downside is the existence of the lowmem kstack mapping, but
    the security guarantees of the feature are preserved.

    In sg_set_buf() (called by sg_init_one and directly) we now check to see if
    the buffer is on the current kernel stack.  If it is, then we redirect the API
    to the lowmem alias of the kernel stack, preserving its assumptions.

    Since the unmapping of the virtual alias can sleep, we need to schedule it
    when called in interrupt context similar to before with vfree.  Unlike before
    however, the contents of the alias depend on the lowmem physical pages, so
    we also need to defer the execution of free_thread_info().

    We also have added a temporary debugging measure for this feature by
    adding a BUG_ON() to virt_to_page() to ensure we're not using a vmapped kernel
    stack address for APIs needing lowmem buffers -- this way we can be notified
    of any other APIs that need similar redirection.

    Thanks to kocka for assisting with some initial qemu/kernel debugging.
```

Thanks for your attention!

# Questions

# References


- https://en.wikibooks.org/wiki/Grsecurity/Application-specific_Settings
- http://pax.grsecurity.net/docs/PaXTeam-H2HC12-PaX-kernel-self-protection.pdf
- https://en.wikibooks.org/wiki/Grsecurity/Appendix/Sysctl_Options
- https://en.wikibooks.org/wiki/Grsecurity/Appendix/Grsecurity_and_PaX_Configuration_Options
- http://pax.grsecurity.net/docs/PaXTeam-LATINOWARE12-PaX-linux-security.pdf
- http://article.gmane.org/gmane.linux.kernel/706950
- http://pax.grsecurity.net/docs/
- https://wiki.gentoo.org/wiki/Hardened/Grsecurity2_Quickstart
- https://wiki.gentoo.org/wiki/Hardened/PaX_Quickstart