

The State of Xtables-addons

Jan Engelhardt

Presented at NFWS 2010

2010-10-18

Table of Contents

1 patch-o-matic

2 Xtables-addons

3 Questions

Table of Contents

1 patch-o-matic

2 Xtables-addons

3 Questions

patch-o-matic (Aug 2002–2007)

- Package to hold **patches** not merged yet in mainline, and those that would never go in anyway.

patch-o-matic (Aug 2002–2007)

- Package to hold **patches** not merged yet in mainline, and those that would never go in anyway.
- A lot of maintenance was involved (50+ patches), each which had to be taken care of when an API change occurred

patch-o-matic (Aug 2002–2007)

- Package to hold **patches** not merged yet in mainline, and those that would never go in anyway.
- A lot of maintenance was involved (50+ patches), each which had to be taken care of when an API change occurred
- Nobody likes that massive sort of maintenance update work

Pitfalls

Worse yet,

- Possibility of incorrect conflict resolution by a novice user.

Pitfalls

Worse yet,

- Possibility of incorrect conflict resolution by a novice user.
- The patch might even apply cleanly and you are just running right into doom.
- Ignoring compiler warnings (a classic) (“it compiles? ship it!”)

Pitfalls

Worse yet,

- Possibility of incorrect conflict resolution by a novice user.
- The patch might even apply cleanly and you are just running right into doom.
- Ignoring compiler warnings (a classic) (“it compiles? ship it!”)
- Stuff you probably never dealt with in your x86-limited world: alignment requirements, endianness

Alignment violation

Unaligned access

```
#define get_u16(X, 0) (*(const __u16 *)((X) + (0)))

if (get_u32(payload, 33) == __constant_htonl(0x71182b1a) &&
    get_u16(payload, 147) == __constant_htonl(0xf792)) {
    printk(KERN_INFO "got WinMX\n");
    return IPP2P_WINMX * 100 + 4;
}
```

- often goes unnoticed because x86 handles it transparently

Resulting oops on sparc64

```
Kernel unaligned access at TPC[79c344] search_winmx+0x123/0x789
```

Safe version

```
#define get_u16(X, 0) get_unaligned((const __u16 *)((X) + (0)))
```

Not good

- silent corruption, kernel oops and an unhappy user.

```

\|/  _ _ _ _ \|/
"@'/ .. \ '@"
/_| \|_/_|_ \|
   \|_U_/_
(sparc64)

```

```

^  ^
--
(xx)\_ _ _ _ _
(_)\      )\ \
U  ||----w |
   ||      ||
(parisc)

```

Not good

- silent corruption, kernel oops and an unhappy user.

<pre> \ / _ _ _ \ / "@'/ .. \ '@" /_ \ _/_ _ \ \ _U_/_ (sparc64) </pre>	<pre> ^ ^ -- (xx)_ _ _ _ _ (_)\)\ \ U ----w (parisc) </pre>
---	--

- did I mention the customers that have to endure a reboot downtime now?

- Why do users think they even need extra modules?
- When was the last time you used `-m fuzzy`?

- Why do users think they even need extra modules?
- When was the last time you used `-m fuzzy`?

Nevertheless, that was p-o-m.

- Don't forget — `ipt_ROUTE` is *gone* — replaced by standard proper policy routing (cf. `iproute2`).

- Why do users think they even need extra modules?
- When was the last time you used `-m fuzzy`?

Nevertheless, that was p-o-m.

- Don't forget — `ipt_ROUTE` is *gone* — replaced by standard proper policy routing (cf. `iproute2`).
- Stop using `ifconfig/route/arp`. Use `iproute2`.

Table of Contents

1 patch-o-matic

2 Xtables-addons

3 Questions

For end users

- No patches, just **module source files**.
- Compile, no reboot needed.
- Much easier for Linux distros to include
- Integrates nicely with kernel updates

For API users

- `compat_xtables.c` acts as an API translator
- Often, `#include "compat_xtables.h"` is all that is needed. Barely any `#ifs` in the extension files.
- Support for a large range of Linux kernel versions, down to Linux 2.6.17 (4 years coverage)

Downsides

- not all features translated to no-ops
 - net namespaces
 - crypto
 - other seldom-used functions
- patching the kernel source, like header files (as IMQ and layer7 require), is not within scope.
 - but you could still make use of the glue code for the parts that do not patch existing files

Submit your module to Xt-a

...or even the kernel, because:

- your code gets an audit
- may even be fixed for you
- less porting work for you by using the compat layer
- benefit from community updates

Documentation:

- “Writing Netfilter modules”
Book in PDF format on <http://jengelh.medozas.de/>

Table of Contents

1 patch-o-matic

2 Xtables-addons

3 Questions

Questions?

- Obligatory URLs
 - `http://xtables-addons.sf.net/`
 - `git://xtables-addons.git.sf.net/gitroot/xtables-addons/xtables-addons`
 - `http://xtables-addons.git.sf.net/` (gitweb)
- Availability
 - Non-exhaustive list 2010: Alpine Linux, Arch, CRUX, Debian, Gentoo, openSUSE, OpenWRT, Polish Linux Distribution (PLD), Slackware
 - `http://freshmeat.net/projects/xtables-addons/`