

You can take the Hacker out
of Perl...

...but you can't take the Perl
out of the hacker.

Laziness, Impatience and Hubris on the Kindle

This talk contains almost no
Perl.

...and almost no Japanese.

I'm very sorry.

And no x86 ASM

I'm not sorry.

The Amazon Kindle

eBook Reader

Showing All 93 Items

By Most Recent First

Spaceman Blues
.....

Brian Francis Slattery

Kindle Download Guide (20...
.....

Feedbooks.com

The New York Times

Mon, Apr 13, 2009

Detroit Metal City v01

Kerouac, Jack - On The Road
.....

◀ Little Brother

Cory Doctorow ▶

The Future of the Internet...
.....

Jonathan Zittrain

Holman2005Science307.1288 (images)

Printcrime
.....

We Haven't Got There Yet
.....

Harry Turtledove

you have the iron self-discipline of a monk.

The good news (for writers) is that this means that ebooks on computers are more likely to be an enticement to buy the printed book (which is, after all, cheap, easily had, and easy to use) than a substitute for it. You can probably read just enough of the book off the screen to realize you want to be reading it on paper.

So ebooks sell print books. Every writer I've heard of who's tried giving away ebooks to promote paper books has come back to do it again. That's the commercial case for doing free ebooks.

Now, onto the artistic case. It's the twenty-first century. Copying stuff is never, ever going to get any harder than it is today (or if it does, it'll be because civilization has collapsed, at which point we'll have other problems). Hard drives aren't going to get bulkier, more expensive, or less capacious. Networks won't get slower or

eInk Screen

800x600; 16 levels of grey

No backlight

Looks great outside!

Good things about the Kindle

3G for 1-click shopping

250,000+ books to buy

(Most < USD10)

Newspaper, Magazine, Blog
subscriptions

Delivered every morning

\$1-\$10 each month

Email → eBook conversion

\$0.10 for autodelivery

\$FREE if you copy the result
by USB

Free Web Browser

Download .mobi, .prc, .azw,
.txt for free

“Experimental”

Means: We might start charging money for this

NetFront 3.5

Basic CSS

Javascript

XmlHttpRequest

Free Wikipedia

Not “experimental”

Text to speech.

Minesweeper!

Find the mines!



Press the M key to mark/unmark mine
Press the R key to restart

M	1	0	0	1	2	M	1
1	1	0	0	1	M	2	1
1	1	1	0	1	1	1	0
1	M	1	0	1	1	1	0
1	1	1	0	1	M	1	0
1	1	0	0	1	1	1	0
M	1	0	0	0	1	1	1
2	2	0	0	1	2	M	1
M	1	1	1	2	M	2	1
1	1	1	M	2	1	1	0

You Won!

Bad things about the Kindle

DRM

When you buy books, they
are locked to the Kindle

Amazon lawyers went after
a tool to let you read non-
Amazon DRMed ebooks

(The same tool can help you
remove the DRM from
books Amazon sells you)

Limited eBook formats

I want to read PDFs

I want to read ePubs

I want to read Manga

Actually, I don't read manga

But lots of my friends do

Surveillance

611 Page

MODEM

Hex ESN: 0x5BA1BD18

Modem Firmware: 131

Slot Cycle Index: 2

Protocol Revision: 6

MSM Version: 6801a

PRL Version: 402

RF Mode: 1-CDMA CELLULAR

Paging Status: IDLE

1xRTT

Phone State: 1-CDMA INIT

SID: 0

NID: 0

Base ID: Not Avail

Latitude: Not Avail

Longitude: Not Avail

BS P REV: 6

Current P REV: 6

Band Class: Not Avail

RF Channel: Not Avail

PN Offset: Not Avail

Ec/lo: Not Avail

RX0 AGC: Not Avail

RX1 AGC: Not Avail

TX Power: Not Avail

TX Gain Adj: Not Avail

TX Power Limit: Not Avail

EVDO

AT State: 1-ACQUISITION

Session State: 0-CLOSED State

Search State: Not Avail

UATI: Not Avail

Color Code: Not Avail

Sector ID: 0x000000

HDRlat: Not Avail

HDRlon: Not Avail

Band Class: Not Avail

EV RF Channel: Not Avail

EV PN Offset: Not Avail

ASET Pilot Energy: Not Avail

RxAGC0: Not Avail

RxAGC1: Not Avail

RxDiv: Not Avail

TxAGC: Not Avail

DRC: Not Avail

SINR: Not Avail

DMD PARAMETERS

Version: mario_1.1.11

Network Signal: (0)No Serv

Bars: 0

RBI: 0

Leap Seconds: 0

Local Offset: 0

MS WAN ACCESS

Data Link: Not Avail

Serving Signal: OUT OF RAN

AN Auth: Not Avail

MIP RRP: Not Avail

Error TS: Not Avail

In the US, 3G includes GPS

The Kindle has a GPS

Device logs are sent to Amazon

Including many user actions

...like the websites you visit

...and what books you read

Appears to include GPS info

Amazon knows where you are

Lock-in

Registration

This device and any content purchased in the Kindle Store are registered to the Amazon.com user name shown below.

Registered User: Jesse R Vincent

[deregister](#)

Registered on Apr 2, 2009

Device Name

The current name for your Kindle is shown below and appears in Home.

Name: Jesse's Kindle

[edit name](#)

Device E-mail

You can send documents to your Kindle's e-mail address shown below. To edit the address or add additional addresses to your approved list of senders, go to www.amazon.com/manageyourkindle.

E-mail: jrv@kindle.com

Personal Info

You can enter personal information such as an address or phone number in case you lose your Kindle.

Personal Information:

[edit personal info](#)

Designed to work only with
Amazon

3G use is free, but only while
Amazon likes you

Summary: The Kindle is an
“appliance”

You wouldn't hack a book,
would you?

I got a Kindle to read books

I didn't plan to hack it

I really wanted to read
books in other formats

I'm a sucker for sexy
platforms

And it was *begging* me

It's a new toy.

I *am* going to hack it.

Laziness

<http://igorsk.blogspot.com>

Hidden debug commands

;debugOn

'help

Search Results: All 82 Items

By Relevance

Little Brother

Cory Doctorow

.....

S

...

K

...

T

...

D

...

K

...

T

...

H

...

P

.....

We Haven't Got There Yet

Harry Turtledove

.....

Private shortcuts: `7777, `allocate,
`applyUpdate, `batteryLoggingDelay,
`checkForUpdate, `compliance,
`consumeMemory,
`countUnmergedDownloadedIndexes,
`disableIndexing, `downloadIndex,
`dumpBattery, `dumpIndexStats,
`einkAdjustments, `help,
`indexForever, `indexStatus, `log611,
`logOpenFiles, `memInfo, `pppStop,
`processTodo, `reloadContentRoster,
`startIndexing, `stopIndexing,
`terminal, `usbNetwork, `usbQa,
`voltLog

close

help

Hubris

“I can’t possibly brick my Kindle with the keyboard, right?”

So, I started typing
commands.

“usbNetwork” sounds
good.

...nothing happened

How about “usbQa”?

It turned off the WIFI

..and turned off USB Disk mode.

Impatience

I gave up

Sometimes laziness wins

An hour later, I rebooted my
Macbook Air



A new network interface has been detected.

The "Ethernet Adaptor (en2)" network interface has not been set up. To set up this interface, open Network Preferences.

Cancel

Network Preferences...

So I set up a DHCP server

Nothing...


```
sh-3.2# tcpdump -i en1  
listening on en1, link-type EN10MB (Ethernet), capture size 96 bytes
```

```
[ ...]
```

```
12:36:15.238229 arp who-has 192.168.15.200 tell 192.168.15.244
```


Network

Show All

Location: Automatic

- AirPort Connected
- Bluetooth Not Connected
- Ethernet Not Connected
- FireWire Not Connected
- Ethern...or (en2) No IP Address

Status: **Unknown State**
The status of your network connection cannot be determined.

Configure: Manually

IP Address: 192.168.15.200

Subnet Mask:

Router:

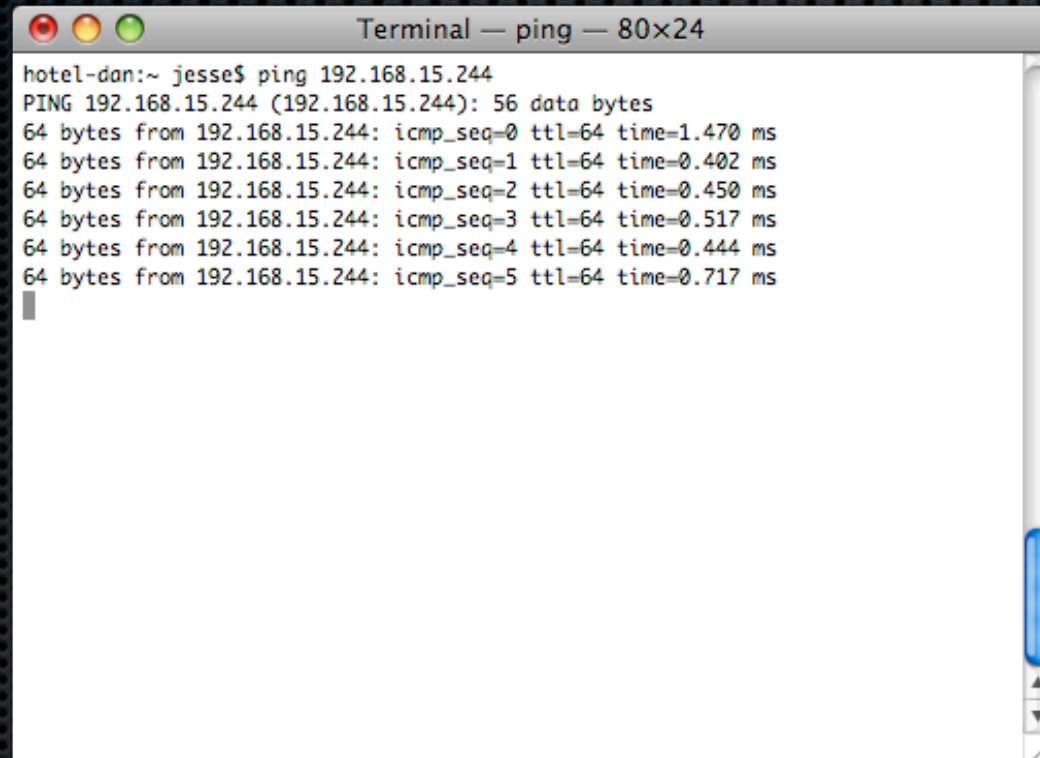
DNS Server:

Search Domains:

Advanced... ?

Click the lock to prevent further changes.

Assist me... Revert Apply

A terminal window titled "Terminal — ping — 80x24" with standard macOS window controls (red, yellow, green buttons). The terminal text shows a successful ping command and its output.

```
hotel-dan:~ jesse$ ping 192.168.15.244
PING 192.168.15.244 (192.168.15.244): 56 data bytes
64 bytes from 192.168.15.244: icmp_seq=0 ttl=64 time=1.470 ms
64 bytes from 192.168.15.244: icmp_seq=1 ttl=64 time=0.402 ms
64 bytes from 192.168.15.244: icmp_seq=2 ttl=64 time=0.450 ms
64 bytes from 192.168.15.244: icmp_seq=3 ttl=64 time=0.517 ms
64 bytes from 192.168.15.244: icmp_seq=4 ttl=64 time=0.444 ms
64 bytes from 192.168.15.244: icmp_seq=5 ttl=64 time=0.717 ms
█
```


Sharing

◀ ▶ Show All 🔍

Computer Name:

Computers on your local network can access your computer at: hotel-dan.local Edit...

On	Service
<input type="checkbox"/>	DVD or CD Sharing
<input type="checkbox"/>	Screen Sharing
<input checked="" type="checkbox"/>	File Sharing
<input type="checkbox"/>	Printer Sharing
<input checked="" type="checkbox"/>	Web Sharing
<input checked="" type="checkbox"/>	Remote Login
<input checked="" type="checkbox"/>	Remote Management
<input type="checkbox"/>	Remote Apple Events
<input type="checkbox"/>	Xgrid Sharing
<input checked="" type="checkbox"/>	Internet Sharing
<input type="checkbox"/>	Bluetooth Sharing

Internet Sharing: On (Sharing your AirPort connection)

Internet Sharing allows other computers to share your connection to the Internet.

Share your connection from:

To computers using:

On	Ports
<input checked="" type="checkbox"/>	Ethernet Adaptor (en2)
<input type="checkbox"/>	FireWire
<input type="checkbox"/>	Ethernet

Click the lock to prevent further changes. ?

Now my Kindle can tether
through my Macbook

I want to read other eBook
formats - attempt #1

The Kindle has a browser

I have a web server

Web based proxy

Small perl app

mobiperl

Perl 4

no strict;

no warnings;

global variables

hmm. no ePub support

Spent a weekend learning
how ePub format works

[github.com/obra/
unsavory-epub-hacks](https://github.com/obra/unsavory-epub-hacks)

Slow. Annoying. Requires a
Server

Servers are evil

Hm.

Back to the drawing board

I want to read other eBook
formats - attempt #2

Let's review what we know:

What's inside the Kindle?

800x600 eInk screen

You know about the screen

Freescale iMX31

ARM1136JF-S

(Includes FPU)

+ Multimedia stuff

2 GB Flash

128 MB RAM

USB OTG + MicroUSB slot

Audio hardware

Keyboard

The Kindle sounds like a
computer, not a book

It **MUST** use some GPL code...

<https://www.amazon.com/gp/help/customer/display.html?inodeId=200203720>

gplrelease.tar.gz

alsa-lib-1.0.13
alsa-lib-1.0.13_patch
alsa-utils-1.0.13
alsa-utils-1.0.13_patch
base-files-3.0.14.ipk
base-passwd_3.5.9
binutils-2.17.50.0.5
bonnie++-1.03c
bootchart-0.9
busybox-1.7.2
dosfstools-2.11
e2fsprogs-1.38
e2fsprogs-1.38_patch
fuse-2.7.1
fuse-2.7.1_link
gcc-4.1.2
glib-2.12.9
glibc-2.5
gst-plugins-base-0.10.17
gst-plugins-base-0.10.6
gststreamer-0.10.17
hotplug-2004_09_20
ifupdown_0.6.8
iptables-1.3.3
klibc-1.5
libol-0.3.18
linux-2.6.22-lab126
lrzsz-0.12.20
lzo-1.08
module-init-tools-3.2.2
module-init-tools-3.2.2_patch
monit-4.9
mtd-utils-1.0.0
picocom-1.4
powertop-1.10
procps-3.2.7
procps-3.2.7_patch
readline-4.3
syslog-ng-1.6.11
sysvinit-2.86
taglib-1.5
uboot-1.3.0-rc3
udev-112
util-linux-2.12r

Linux

I can work with this

But how do I get code onto it?

My friend nmap tells me...

The Kindle listens on a few
ports.

None of them love me at all

I guess I'll need to take
matters into my own hands.

Hey, the Kindle has busybox

busybox has telnetd

Maybe I just need to install
`/etc/rc5.d/S99telnetd`

More research from
<http://igorsk.blogspot.com>

Kindle 1 update extractor

(Python script)

But I want to make new
updates...

I reverse engineered the
reverse engineering tool

Updates are a short header,
an MD5 and a tarball

...run through a trivial cipher

they're **not** encrypted

The Kindle2 is a little
different than the Kindle1

It has different magic #s
in the update header

I waited for the first Kindle 2
“update.bin”

I grabbed its header

I built a new “update”

It installed one file

/etc/rc5.d/S99telnetd


```
/bin/busybox telnetd -p 2323
```


...nope.

I ran strings on the Kindle's
busybox

No telnetd!

Where can I get a busybox
for the Kindle?

What else has a a similar
CPU?

My gPhone!



Lots of people built static
busybox for the gPhone

telnetd: take 2

login:

login: root

Password:

Login incorrect



`/bin/sh`

makes a better

`/bin/login`




```
125-6-81-160:ß jesse$ telnet kindle 2323
Trying 192.168.15.244...
Connected to kindle.
Escape character is '^Ü'.
```

```
/ # cat /etc/motd
```

```
#####
```

```
# NOTICE * NOTICE * NOTICE #
```

```
#####
```

```
Rootfs is mounted read-only. Invoke mntroot rw to
switch back to a writable rootfs.
```

```
#####
```

```
/ #
```


What I found:

Most of /sbin is written in sh

Fun stuff in /proc

/proc/config.gz


```
#  
# Automatically generated make config: don't edit  
# Linux kernel version: 2.6.22.19  
# Mon Mar  2 12:13:07 2009  
#  
CONFIG_ARM=y  
CONFIG_SYS_SUPPORTS_APM_EMULATION=y  
# CONFIG_GENERIC_GPIO is not set  
CONFIG_GENERIC_TIME=y  
CONFIG_GENERIC_CLOCKEVENTS=y  
CONFIG_MMU=y  
# CONFIG_NO_IOPORT is not set  
CONFIG_GENERIC_HARDIRQS=y  
CONFIG_STACKTRACE_SUPPORT=y  
CONFIG_LOCKDEP_SUPPORT=y  
...
```


...I could rebuild the Kernel

/proc/filesystems

nodev sysfs
nodev rootfs
nodev bdev
nodev proc
nodev sockfs
nodev pipefs
nodev
anon_inodefs
nodev futexfs
nodev tmpfs
nodev
inotifyfs

nodev devpts
nodev ext3
nodev ramfs
nodev msdos
nodev vfat
nodev nfs
nodev
rpc_pipefs
nodev fuse
nodev fuseblk
nodev fusectl



I'm not restricted to 2GB

It's a Linux box

I can cross-compile!

<http://www.codesourcery.com/>

Prebuilt ARM toolchain

It generates generic ARM
machine code

That's ok, but not great

I'll cross compile Perl

1 day of frustration passes

I won't cross-compile Perl

Crosscompiling Perl

==

Bad Joke

I'll cross-compile Python

Same bad joke

Maybe I need a native
compiler for ARM

Where do I get an ARM
build farm?

I have a gPhone



It's not a great build host

It's not a great build host...

...but it has an important
advantage


```
apt-get install gcc
```


N810: Linux 2.6; glibc 2.5

N810 binaries run
unmodified on the Kindle

I built perl in an hour

Sadly, I realized that Python
is a better choice

I also realized that building on the Kindle works better than on the N810.

(Version skew in extra
libraries makes things hard)

I tried building gcc on the
N810...

Found Pengutronix /
OSELAS.de

It's a compiler toolchain
builder.

I built my own crosscompilers
for ARM1136JF-S - Linux 2.6
- glibc 2.5

I used the cross compiler to compile gcc, glibc (for proper headers), binutils, shellutils, dropbear & screen

| cross-compiled nfsmount

I nfs-mounted a disk image
with the compiler

Then I started building more
stuff

I am a Perl Hacker

I believe in the three virtues

Lazyness

Impatience

Hubris

Sometimes, Perl isn't the
Right Tool

Calibre is the Killer App for ebook conversion and management

Library
Reader
191 MB available

Search: Search (For Advanced Search click the button to the left)

	Title	Author(s)	Size (MB)	Date	Rating	Publisher	Tags
25	Saturday	Ian McEwan	0.5	09 Mar 2008		London : Jonathan Cape, 2005.	
26	Knife of Dreams	Robert Jordan	0.9	07 Mar 2008			
27	Hunter's Run	George R. R. Martin Gardner Dozois Daniel Abraham	0.4	04 Mar 2008	★ ★ ★	Eos	scifi

Hunter's Run
george r. r. martin
gardner dozois
daniel abraham

HUNTER'S RUN

Formats: lrf, rar

Comments: SUMMARY: Like so many others, Ramón Espejo ran from the poverty and hopelessness of the Third World to the promise of a new world—joining a host of like-minded workers and dreamers aboard one of the great starships of the mysterious, repulsive Erye. But the life he found on the far-off planet of São Paulo was no better than the one he had abandoned. Tough, volatile, and angry—a luckless emigrant hoping for that one rich strike that will make him wealthy—Ramón is

Lazyness

I like the best tools

The best tools already exist

Calibre has dozens of eBook
format converters.

Why reimplement them?

Hubris

“I can learn enough Python
in a weekend to port this
application to the Kindle”

The downside

Dependencies

Who's dealt with Python app dependencies?

No CPAN.

Everything you need is in the
Standard Library.

If it's not in the Standard
Library, it's not worth using.

Except when you need it.

They have....

“easy_instal”

It's not so easy

It *is* very perlish

It does recursive web
scraping to find tarballs on
developers' web sites.

Most of the deps actually
installed ok.

I just ran the app over and over until it stopped erroring.

God I miss Perl.

And then we get to the big
problem.

Qt

Calibre's UI is in Qt

...so its backend uses Qt
because it's easy

PyQt binds Qt to Python

For Qt for Windows

for Qt for Mac

for Qt for X11

No X11 on Kindle

(Just a Framebuffer)

QtEmbedded

No problem!

No PyQtEmbedded

Finally got Calibre running...

by hacking out components
I don't need.

It was good enough to try to
convert a trivial ebook.

It took 12 hours...

...after I built swaputils and
gave it 256MB of swap



So what was it doing?

HTML → Mobipocket
converter

With a full CSS engine

It visits every DOM
element...

and computes CSS styles to
convert them to trivial HTML
3.2...

...twice.

Lazyness, Impatience,
Hubris can all help here.

Help me Larry-wan.

Very few CSS rules really matter.

The Kindle supports very
little HTML.

It mostly supports HTML
3.2...just no `<pre>`

...that's the only thing the 12
hour CSS engine got us

You can emulate `<pre>` with
`<tt>` and ` `;

EVERYBODY STAND BACK.



I KNOW REGULAR EXPRESSIONS.




```
if tag == 'pre':  
    self.inside_pre = 1  
    tag = 'tt'
```

```
if prefixname(elem.tag, nsrmap) == 'pre':  
    buffer.write('<br/>\n')  
    self.inside_pre = 0
```

```
if self.inside_pre:  
    text=text.replace(' ', '&nbsp;')  
    text=re.sub(r'(\r\n|\r|\n)', '<br/>\n', text)
```


Now it runs in 60 megs and
about 10 minutes

So, now I can run code.

Still no UI access.

I don't really want to hack
Java GUI code.

And where could I plug my custom UI into the Kindle's?

I don't want to break
Amazon's UI.

Oh hey.

There *is* an application I
could replace with
something custom...

Find the mines!



Press the M key to mark/unmark mine
Press the R key to restart

M	1	0	0	1	2	M	1
1	1	0	0	1	M	2	1
1	1	1	0	1	1	1	0
1	M	1	0	1	1	1	0
1	1	1	0	1	M	1	0
1	1	0	0	1	1	1	0
M	1	0	0	0	1	1	1
2	2	0	0	1	2	M	1
M	1	1	1	2	M	2	1
1	1	1	M	2	1	1	0

You Won!

But really, I don't want to.

Sure, I could decompile.

It's obfuscated.

It'd be annoying.

If I built UI, I'd have to
maintain a UI.

And users can break a UI.

No buttons

=


Less to screw up

But I have this ebook
converter.

I do want to let users
convert books.

What to do?

nodev	sysfs	nodev	devpts
nodev	rootfs		ext3
nodev	bdev	nodev	ramfs
nodev	proc		msdos
nodev	sockfs		vfat
nodev	pipefs	nodev	nfs
nodev		nodev	
anon_inodefs		rpc_pipefs	
nodev	futexfs	nodev	fuse
nodev	tmpfs		fuseblk
nodev		nodev	fusectl
inotifyfs			



Inotify blocks on filesystem
events.

pyInotify lets me get at fs
events easily.


```
class InotifyListener (threading.Thread):
    global cv
    def run ( self ):
        global conversionQueue

        wm = WatchManager() # Watch Manager
        mask = IN_MOVED_TO | IN_CREATE # watched events

        p = PTmp()
        notifier = Notifier(wm, p)
        wdd = wm.add_watch('/mnt/us/documents', mask, rec=True)
        notifier.loop()
```


It works great for downloads

Copies over USB didn't
trigger inotify events.

It's probably something
fuse-related.

I went for the cheap hack.

When you eject the Kindle, it generates a DBus event.


```
class DbusWatcher (threading.Thread):
    global cv
    def run ( self ):
        global conversionQueue
        cmd='/usr/bin/dbus-monitor --system'
        pipe = subprocess.Popen(cmd, shell=True,
stdout=subprocess.PIPE).stdout
        while 1:
            line = pipe.readline()
            if any(line.find(i) != -1 for i in ['usbPlugOut', 'resuming']):
                for f in os.listdir('/mnt/us/documents'):
                    maybe_enqueue_file('/mnt/us/documents/'+f)
```


What's next?

Remember config.gz?

I can build a new kernel

...and add back missing
drivers

USB Mass Storage Host

USB WIFI?

What isn't next?

Reverse engineering Java to
extend the Kindle's UI

Python, and Shell, I'm happy
to hack for a good cause.

Java is another matter
entirely.

Thanks!

I had a big finish planned.


I was going to build and
show off a manga converter.

(for .cbz format books)

So I downloaded a .cbz.

...and copied it to the Kindle...

...and I saw this...



The best hacking
is no hacking.

Thanks!