

A Large Scale Analysis of the Security of Embedded Firmware

A. Costin, J. Zaddach, A. Francillon, D. Balzarotti
EURECOM, France

SECURE 2014, Warsaw

Who are we?




Aurélien Francillon



Davide Balzarotti




Andrei Costin



Jonas Zaddach

Embedded Systems Are Everywhere

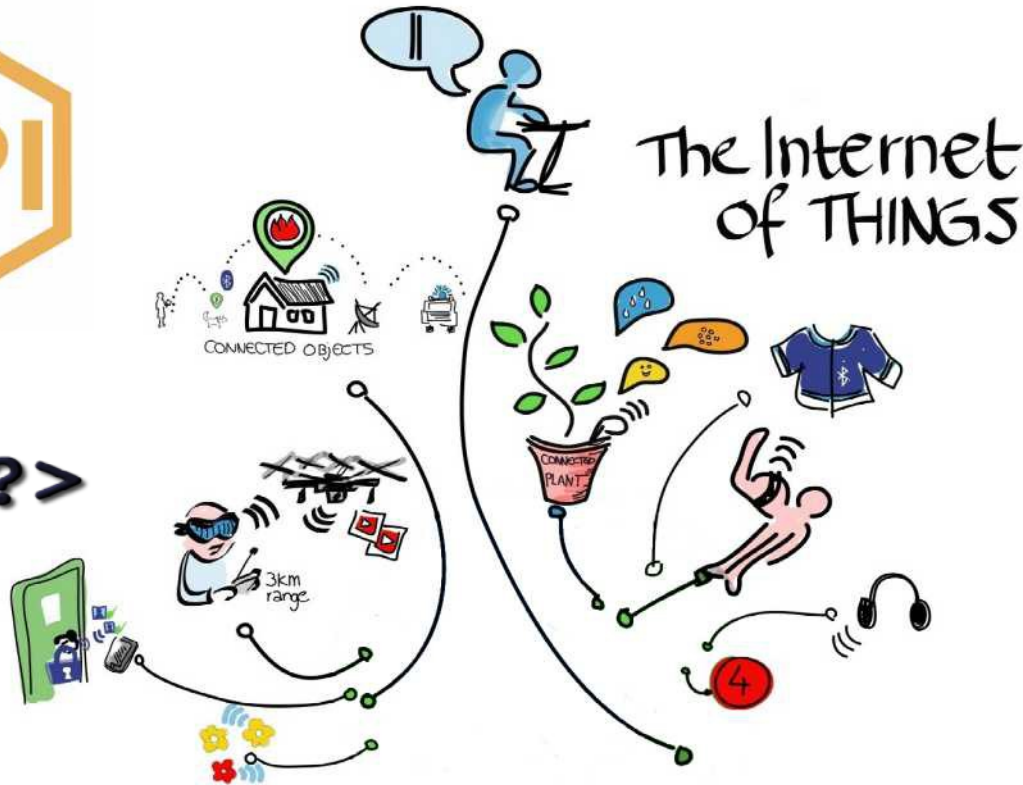


by Wilgebroed on Flickr [CC-BY-2.0]

Smarter & More Complex



<?xml?>



{JSON}
RESTful API
GET PUT POST DELETE

CONNECT

by Wilgenbroed on Flickr [CC-BY-2.0]

Interconnected



<?xml?>



CONNECT THE WORLD

by Wilgenbroed on Flickr [CC-BY-2.0]

{JSON}

RESTful API
GET PUT POST DELETE



Many Examples of Insecure Embedded Systems

- Routers



Firefox Reverse Engineering a D-Link B...
www.devttys0.com/2013/10/reverse-engineering-a-d-link-backdoor/

Based on the source code of the HTML pages and some Shodan [search results](#)
D-Link devices are likely affected:

- DIR-100
- DIR-120
- DI-624S
- DI-524UP
- DI-604S
- DI-604UP
- DI-604+
- TM-G5240

Additionally, several Planex routers also appear to use the same firmware:


- BRL-04R
- BRL-04UR
- BRL-04CW

You stay classy, D-Link.

Many Examples of Insecure Embedded Systems

- Routers
- Printers

Networked printers at risk
(30/12/2011, McAfee Labs)



Many Examples of Insecure Embedded Systems

- Routers
- Printers
- VoIP

Cisco VoIP Phones Affected By On Hook Security Vulnerability (12/06/2012, Forbes)



Many Examples of Insecure Embedded Systems

- Routers
- Printers
- VoIP
- Cars

Hackers Reveal Nasty New Car Attacks – With Me Behind The Wheel (12/08/2013, Forbes)




Many Examples of Insecure Embedded Systems

- Routers
- Printers
- VoIP
- Cars
- Drones

Hacker Releases Software to Hijack Commercial Drones


by BRYANT JORDAN on DECEMBER 9, 2013

 Like 489 people like this. Be the first of your friends.



Many Examples of Insecure Embedded Systems

- Routers
- Printers
- VoIP
- Cars
- Drones
- ...



The Goal

Perform a large scale analysis
to gain a better understanding
of firmware problems



The Problem With Large Scale Analysis

- Heterogeneity of
 - Hardware, architectures, OSes
 - Users, requirements
 - Security goals

The Problem With Large Scale Analysis

- Heterogeneity of
 - Hardware, architectures, OSes
 - Users, requirements
 - Security goals
- Manual analysis does not scale, it requires
 - Finding and downloading firmware
 - Unpacking and initial analysis
 - **Re**-discovering a similar bugs

Previous Approaches

- Test on real devices [Bojinov09CCS]
 - Accurate results
 - Does not scale well

Previous Approaches

- Test on real devices [Bojinov09CCS]
 - Accurate results
 - Does not scale well
- Scan devices on the Internet
 - Large scale testing [Cui10ACSAC]
 - Can only test for known vulnerabilities
 - Blackbox approach
 - More is too intrusive [Census2012]

Our Approach to The Large Scale Analysis

- Collect a large number of firmware images

Our Approach to The Large Scale Analysis

- Collect a large number of firmware images
- Perform broad but simple static analysis

Our Approach to The Large Scale Analysis

- Collect a large number of firmware images
- Perform broad but simple static analysis
- Correlate across firmwares

Our Approach to The Large Scale Analysis

- Collect a large number of firmware images
- Perform broad but simple static analysis
- Correlate across firmwares
- Advantages
 - No intrusive online testing, no devices involved
 - Scalable

Our Approach to The Large Scale Analysis

- Collect a large number of firmware images
- Perform broad but simple static analysis
- Correlate across firmwares
- Advantages
 - No intrusive online testing, no devices involved
 - Scalable
- Many challenges remain

Mainstream Systems Have Centralized Updates

The image displays two software update interfaces side-by-side. On the left is the Ubuntu Update Manager, and on the right is the Windows Update control panel. A central dialog box from an Apple system is overlaid on the Windows Update window.

Ubuntu Update Manager:

- Welcome to Ubuntu
- Since this version of Ubuntu was released, these software updates have been issued.
- Distribution updates**
- Daemon which notifies about package updates (Size: 53 kB)
- Files shared between update-notifier and update-notifier-common (Size: 48 kB)
- 2 updates have been selected. 100 kB will be downloaded.
- Settings...

Windows Update:

- Control Panel Home
- Windows Update
- Search Control Panel
- File Edit View Tools Help
- Install updates for your computer
- 1 important update is available
- 8 optional updates are available
- 1 important update selected, 7.4 MB
- Install updates
- Most recent check for updates: Today at 8:12 AM
- Updates were installed: 3/15/2010 at 4:05 PM
- View update history
- Next time I want to receive updates: For Windows and other products from Microsoft Update
- Find out more about free software from Microsoft Update. Click here for details.

Apple Software Update Dialog:

Software Update

New software is available for your computer.
If you're not ready to install now, you can use the Software Update preference to check for updates later.

Install	Name	Version	Size
<input checked="" type="checkbox"/>	iDVD Update	7.0.1	20.2 MB
<input checked="" type="checkbox"/>	iMovie Update	7.1	46.1 MB
<input checked="" type="checkbox"/>	GarageBand Update	4.1	47.0 MB
<input checked="" type="checkbox"/>	iPhoto Update	7.1	61.8 MB
<input checked="" type="checkbox"/>	iLife Support	8.1	10.3 MB

This update supports system software components shared by all iLife '08 applications, improves overall stability, addresses a number of other minor issues, and supports general compatibility issues. It is recommended for all users of iLife '08.

Note: Use of this software is subject to the original Software License Agreement(s) that accompanied the software being updated. A list of Apple SLAs may be found here: <http://www.apple.com/legal/sla/>.

Quit Install 5 Items

Challenge: Embedded Systems

Update Sources are diverse

- Public site
 - Manufacturer web site
 - FTP site
- Hidden site
 - Accessed by firmware update utility
- Restricted site
- Request-only updates
- Delivery on other media (CD-Rom, ...)
- Firmware only delivered on device

Challenge: Embedded Systems Update Mechanisms are diverse



Collecting a Dataset

- No large scale firmware dataset yet
 - As opposed to existing datasets in security or other research areas

Collecting a Dataset

- No large scale firmware dataset yet
 - As opposed to existing datasets in security or other research areas
- We collected a subset of the firmwares available for download

Collecting a Dataset


- No large scale firmware dataset yet
 - As opposed to existing datasets in security or other research areas
- We collected a subset of the firmwares available for download
- Still many firmwares are not publicly available

Collecting a Dataset

- No large scale firmware dataset yet
 - As opposed to existing datasets in security or other research areas
- We collected a subset of the firmwares available for download
- Still many firmwares are not publicly available
 - www.firmware.re project

Challenge: Firmware Identification


← Clearly a Firmware



Challenge: Firmware Identification

← Clearly a Firmware


Clearly not a Firmware →



Challenge: Firmware Identification

← Clearly a Firmware

Clearly not a Firmware →



Challenge: Firmware Identification

- E.g., upgrade by printing a PS document

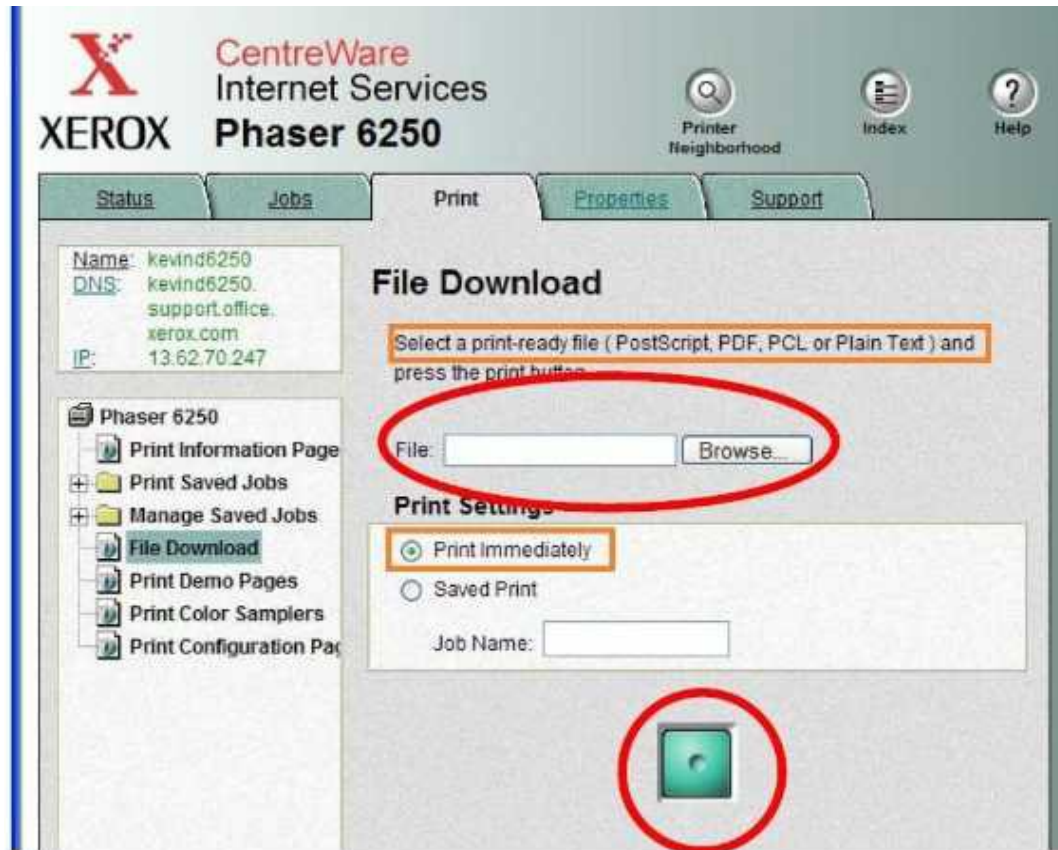


Figure 4: Select the firmware update file and press the green button to send it.

Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?

Challenge: Unpacking & Custom Formats

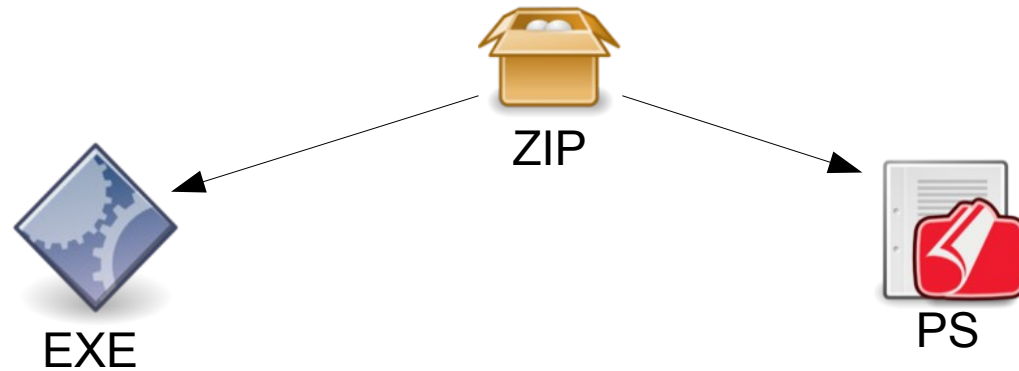
- How to reliably unpack and learn formats?



ZIP

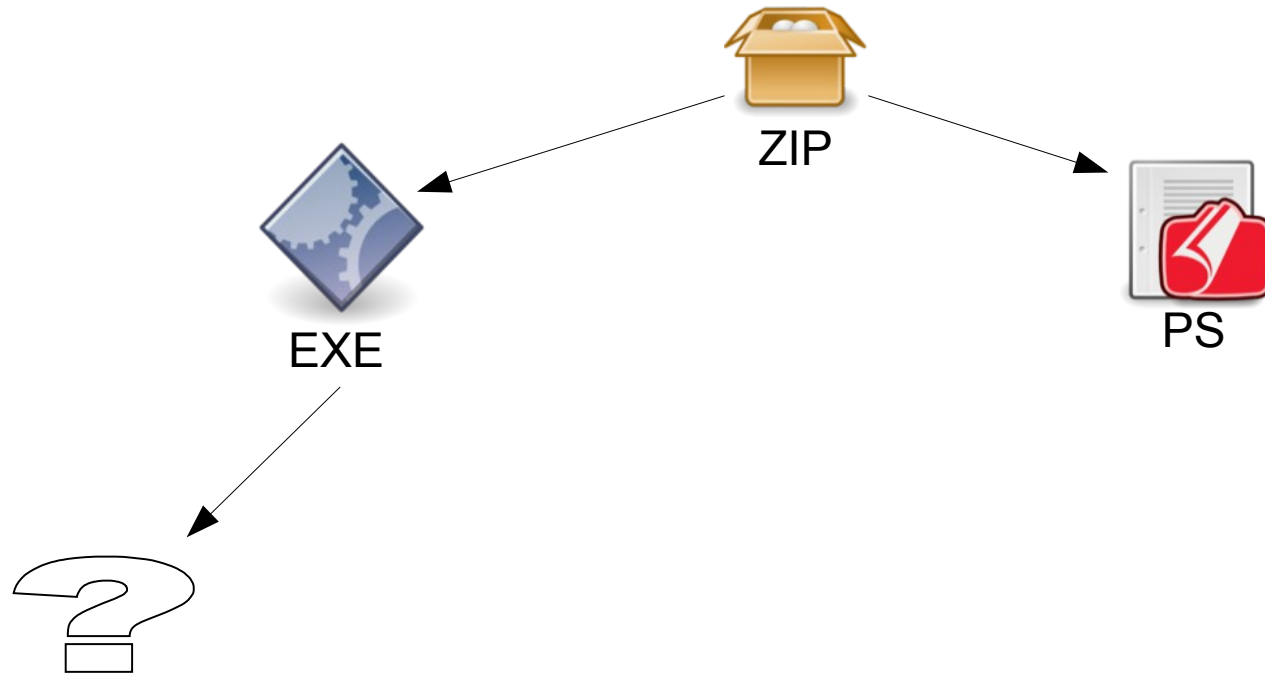
Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?




Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?



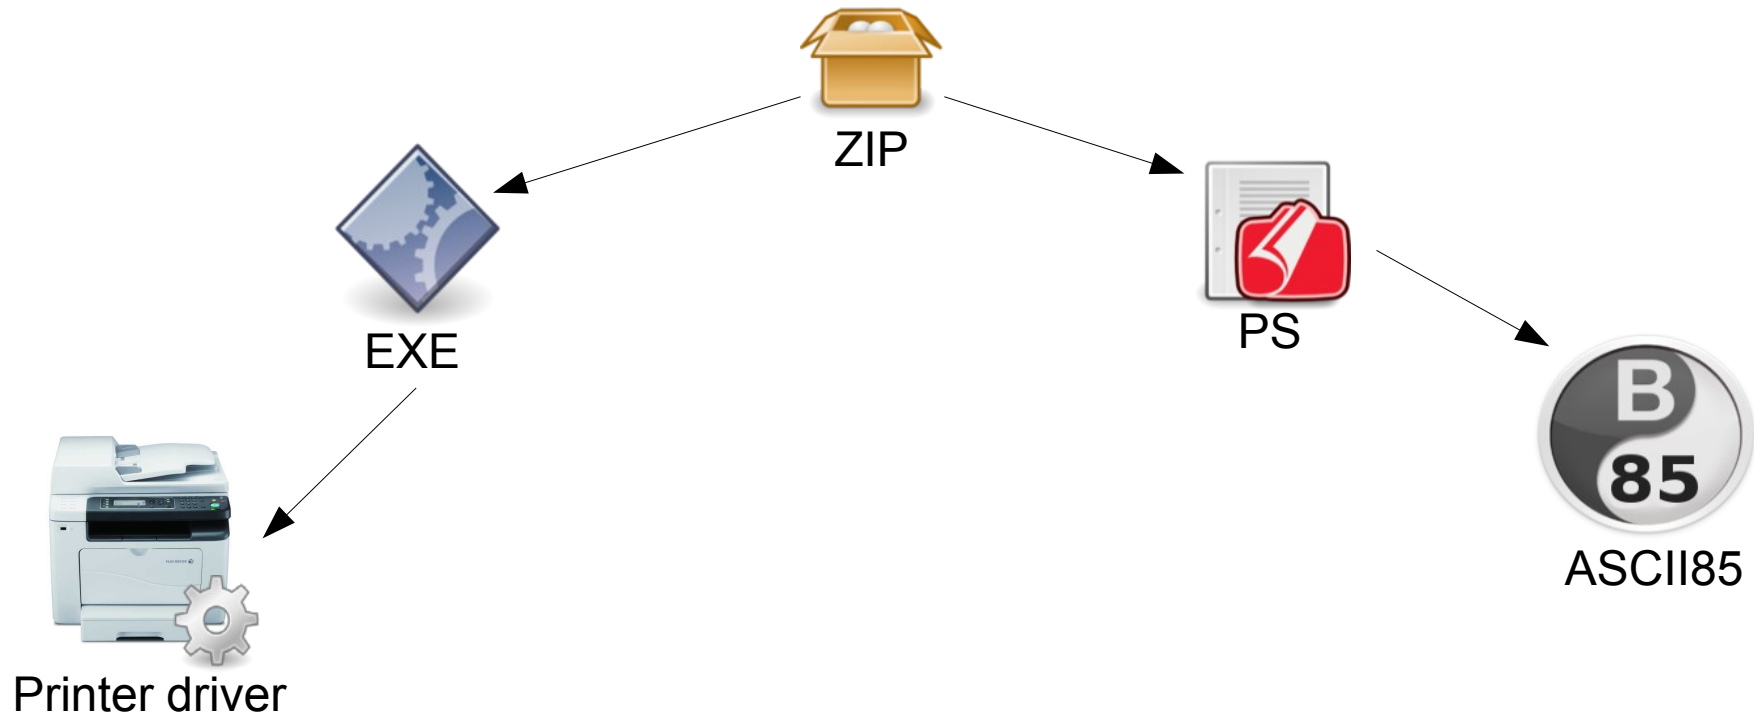
Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?



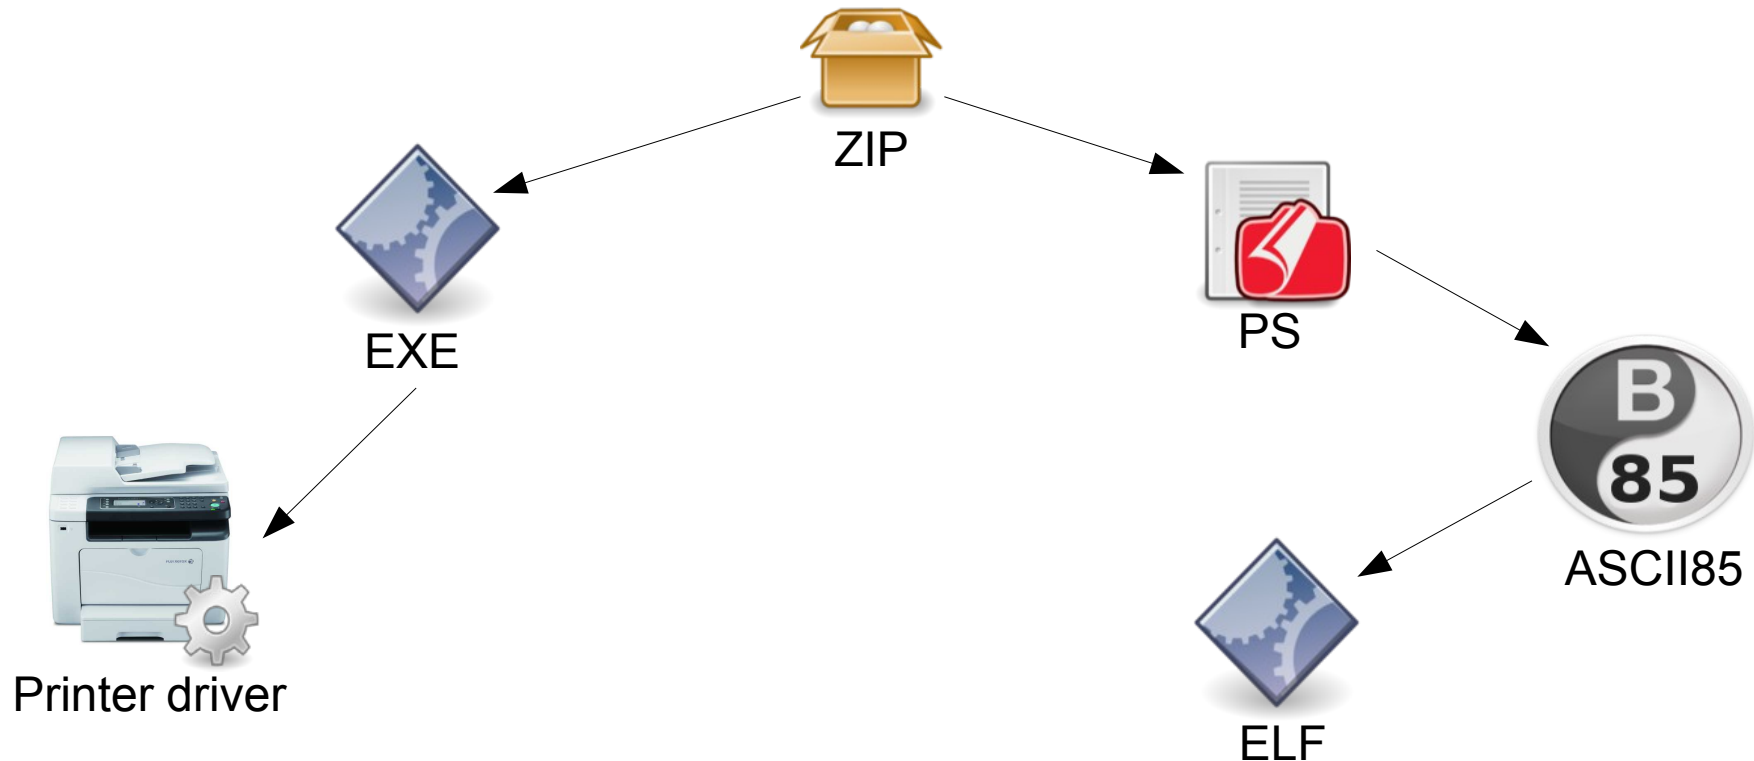
Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?



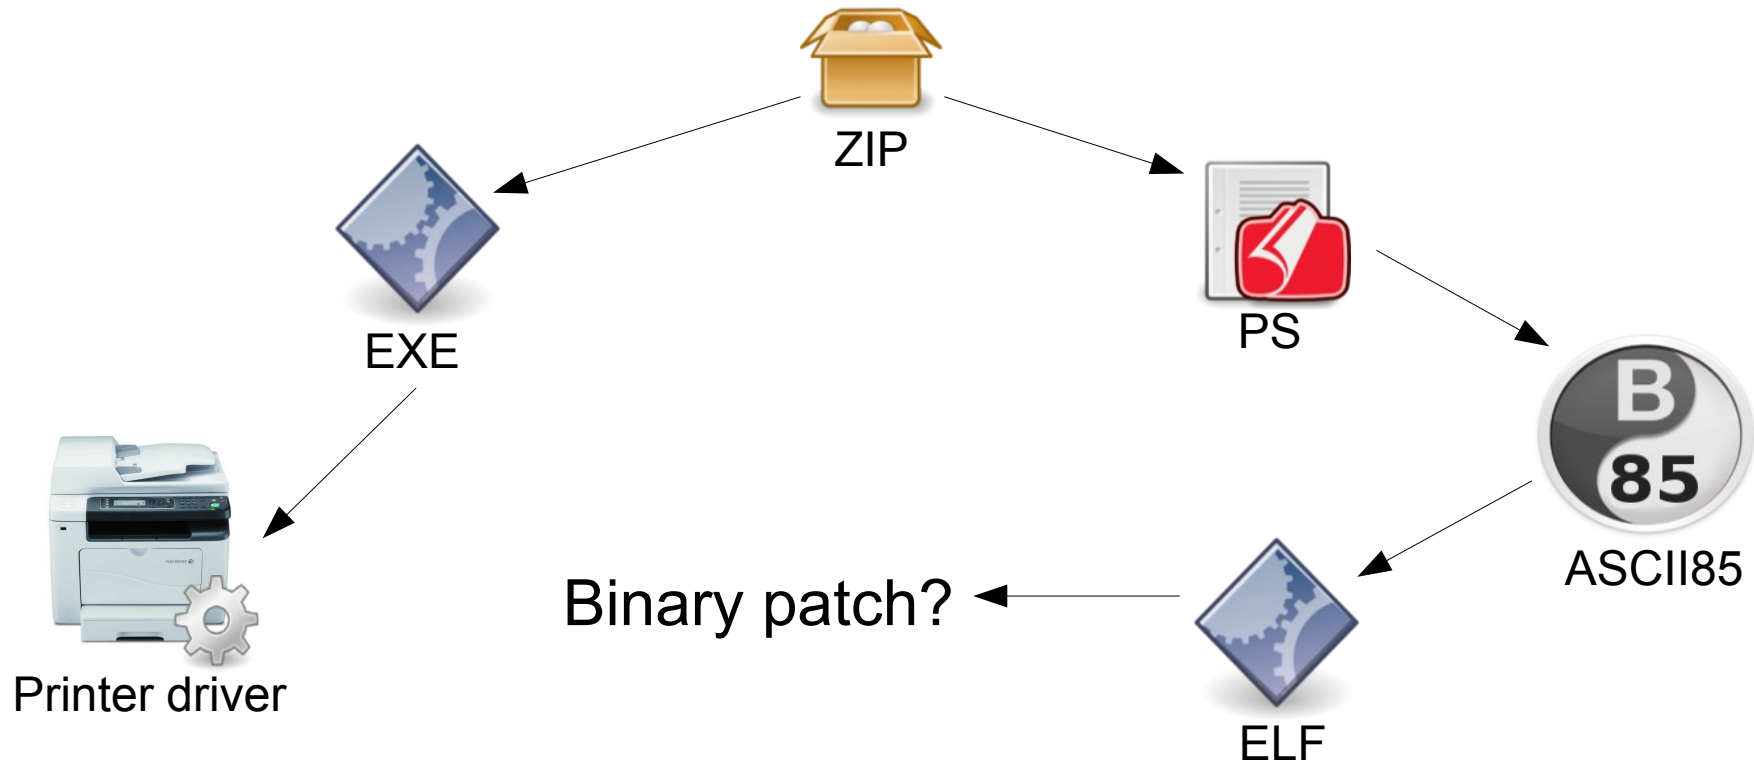
Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?



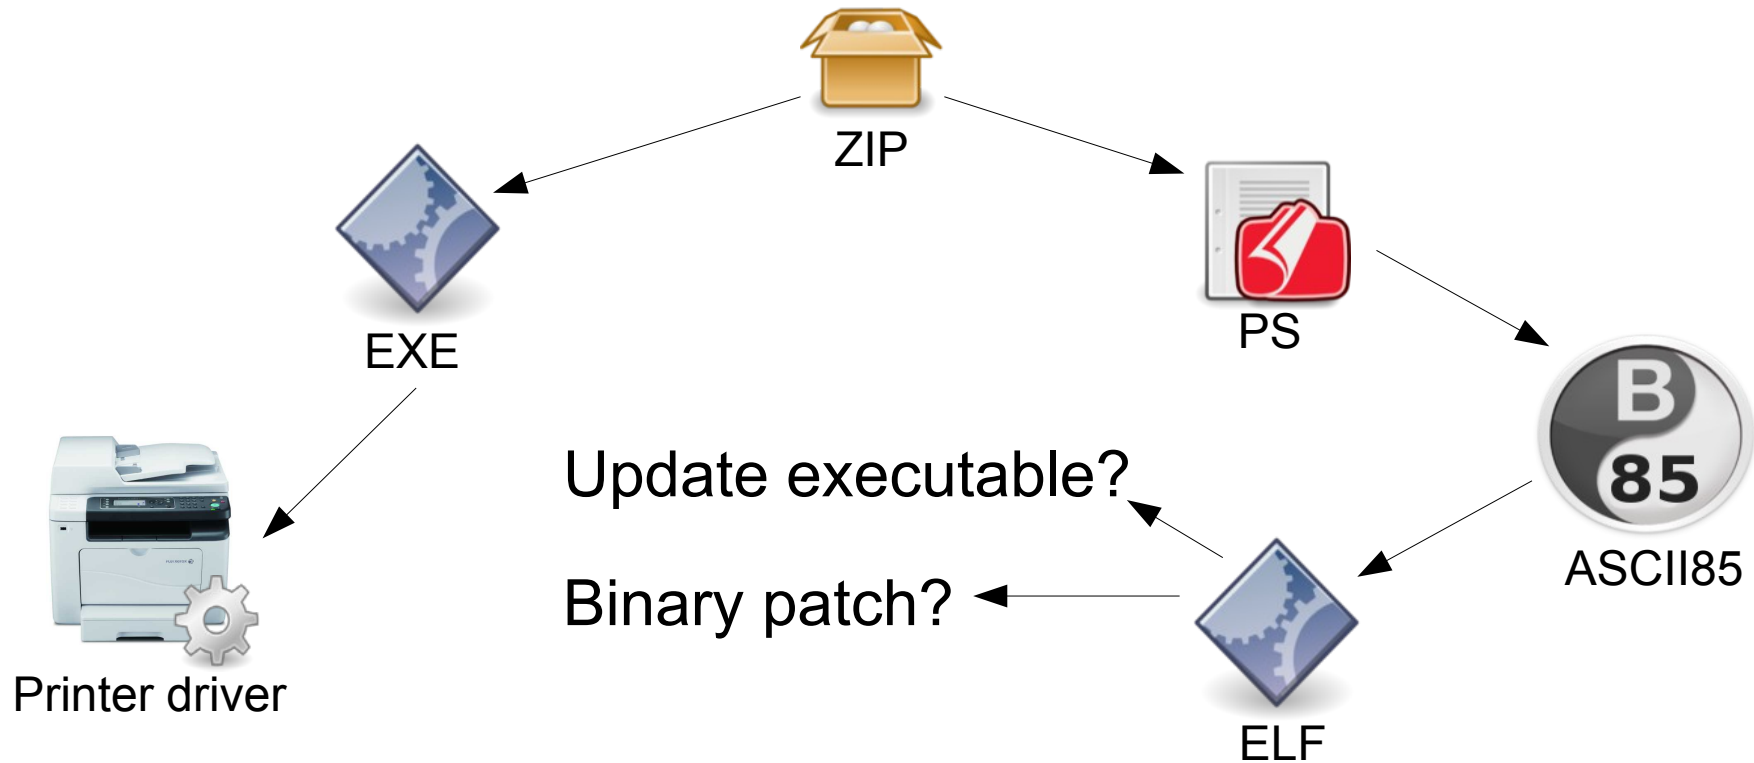
Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?




Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?



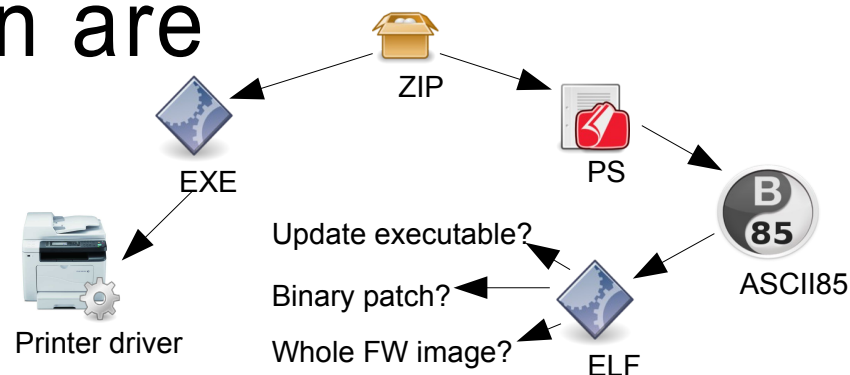
Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?



Challenge: Unpacking & Custom Formats

- How to reliably unpack and learn formats?
- Firmware updates often are "russian dolls"
- Sometimes result of unpacking is just a binary data blob



Our Approach to Unpacking & Custom Formats

- Often a firmware image is just a binary blob
 - File carving required
 - Bruteforce at every offset with all known unpackers
 - Have good heuristics when to stop carving

Our Approach to Unpacking & Custom Formats

- Often a firmware image is just a binary blob
 - File carving required
 - Bruteforce at every offset with all known unpackers
 - Have good heuristics when to stop carving
- We compared existing tools and used BAT (Binary Analysis Toolkit)
 - Supports recursive extraction and carving
 - Extended it with multiple custom unpackers

Challenge:

Scalability & Computational Limits

- Unpacking and file carving is very CPU intensive


Challenge:

Scalability & Computational Limits

- Unpacking and file carving is very CPU intensive
- Results in millions of unpacked files
 - Manual analysis infeasible
 - One-to-one fuzzy hash comparison is CPU intensive

Challenge: Scalability & Computational Limits

- Fuzzy hashing becomes difficult with lots of files



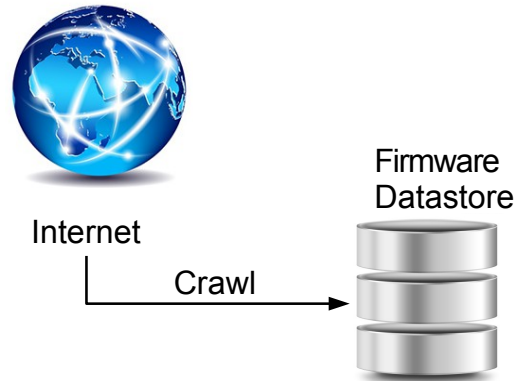
Challenge: Results Confirmation

- An issue found statically
 - Cannot guarantee exploitability
 - May not apply to a real device
 - E.g., vulnerable daemon present but never started

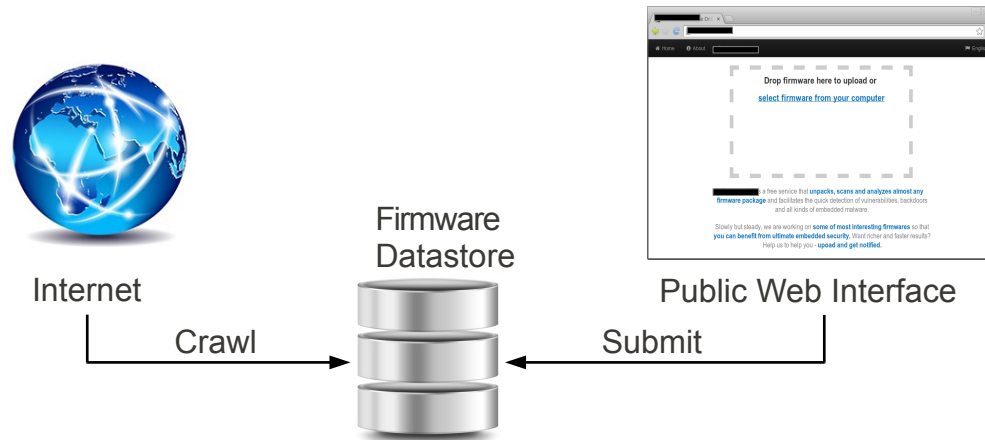
Challenge: Results Confirmation

- An issue found statically
 - Cannot guarantee exploitability
 - May not apply to a real device
 - E.g., vulnerable daemon present but never started
- Issue confirmation is difficult
 - Requires advanced analysis (static & dynamic)
 - Does not scale for heterogeneous firmware
 - Often requires real embedded devices

Architecture



Architecture




Architecture



Internet

Crawl

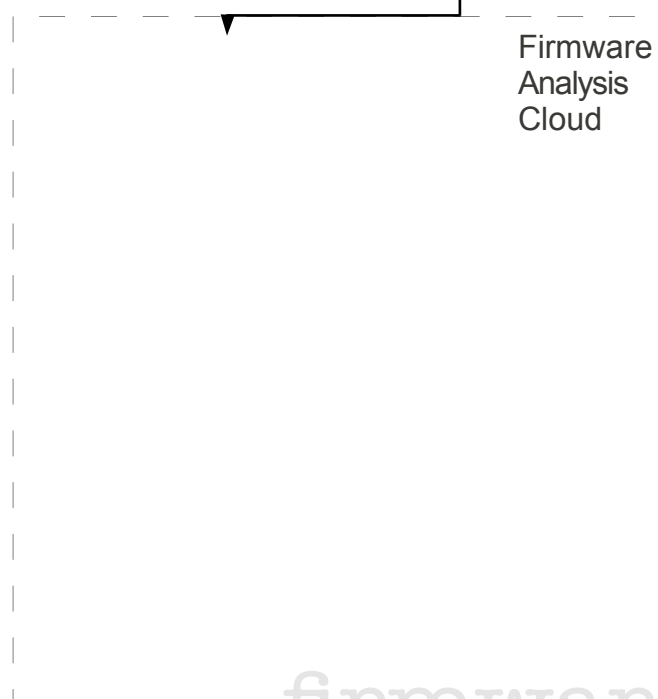
Firmware
Datastore



Public Web Interface

Submit

Firmware
Analysis
Cloud



firmware · əɪ


Architecture



Internet

Crawl

Firmware
Datastore



Public Web Interface

Submit

Firmware
Analysis
Cloud



Master

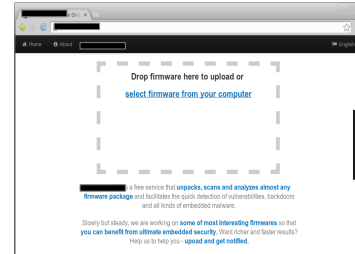

Architecture



Internet

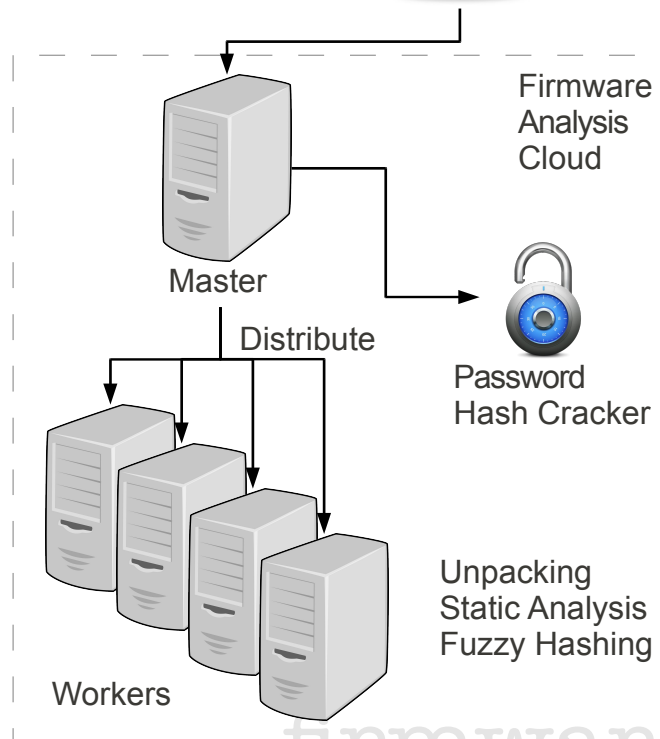
Crawl

Firmware
Datastore

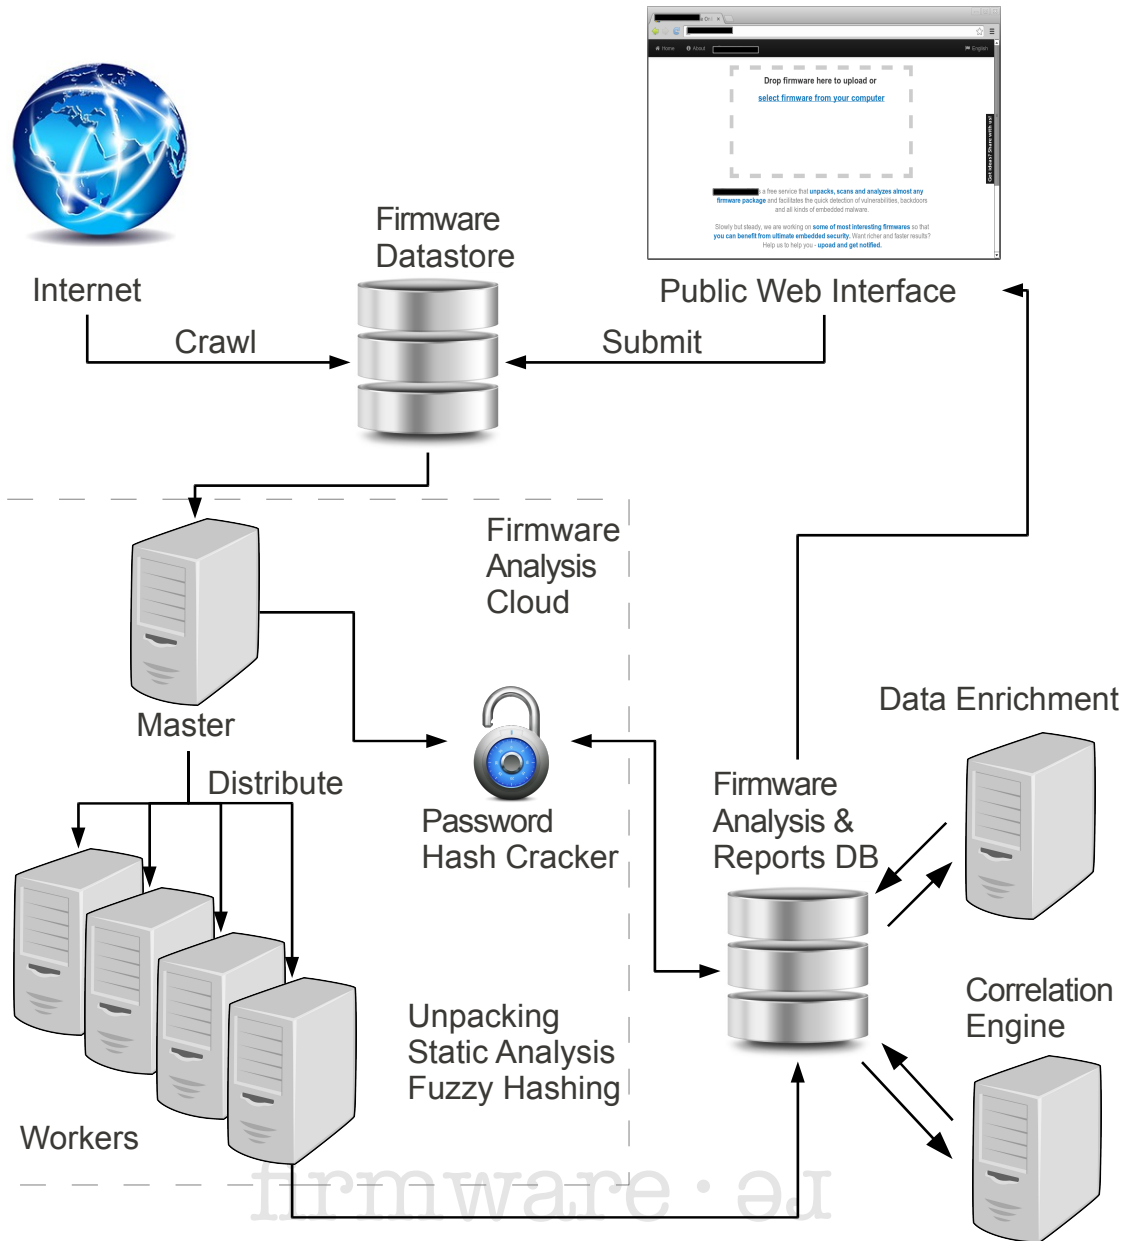


Public Web Interface

Submit



Architecture



Crawler

- Multiple seeds
 - FTP-index engines
 - Google Custom search engines
- Several download techniques
 - WGET scripts
 - Beautiful Soup scripts
- 759 K collected files, 1.8 TB of disk space

www.Firmware.RE (beta)

Will provide Unpacking and Analysis

The screenshot shows a web browser window with the address bar displaying 'www.firmware.re'. The page title is 'firmware · 01 (beta)'. The main content area features a blue 'Upload' button, followed by 'Info' and 'Examples' links. A large dashed box contains the text: 'To start, drag-n-drop firmware here or [select firmware from your computer](#)'. The footer includes links for 'Blog', 'Twitter', 'contact@firmware.re', 'Google groups', 'ToS', and 'Privacy policy'. A vertical banner on the right side reads 'Got ideas? Share with us!'.

www.Firmware.RE (beta)

Will provide Unpacking and Analysis

Firefox | firmware · ə - Free Online Firm... | +

www.firmware.re | Google | Home | +

firmware · ə (beta) | USENIX Security '14 | BH13US | About | English

Upload
Info
Examples

To start, drag-n-drop firmware here or [select firmware from your computer](#)


IMPROVED

Got ideas? Share with us!

Blog | Twitter | contact@firmware.re | Google groups | ToS | Privacy policy


Unpacking

- 759 K total files collected




Unpacking

- 759 K total files collected
- ↓ Filter non firmware
- 172 K filtered interesting files




Unpacking

- 759 K total files collected
 - ↓ Filter non firmware
- 172 K filtered interesting files
 - ↓ Random selection
- 32 K analyzed



Unpacking

- 759 K total files collected
 - ↓ Filter non firmware
- 172 K filtered interesting files
 - ↓ Random selection
- 32 K analyzed
 - ↓ Successful unpack
- 26 K unpacked (fully or partially)



Unpacking


- 759 K total files collected
 - ↓ Filter non firmware
- 172 K filtered interesting files
 - ↓ Random selection
- 32 K analyzed
 - ↓ Successful unpack
- 26 K unpacked (fully or partially)
 - ↓ Unpacked files
- 1.7 M resulted files after unpacking

Static Analysis

- Misconfigurations
 - Web-server configs, Credentials, Code repositories
- Data enrichment
 - Version banners → Software packages and versions
 - Keywords → Known problems (e.g., telnet, shell, UART, backdoor)
- Correlation/clustering
 - Fuzzy hashes, SSL certificates, Credentials


Example: Correlation

- Correlation via fuzzy-hashes (ssdeep, sdhash)



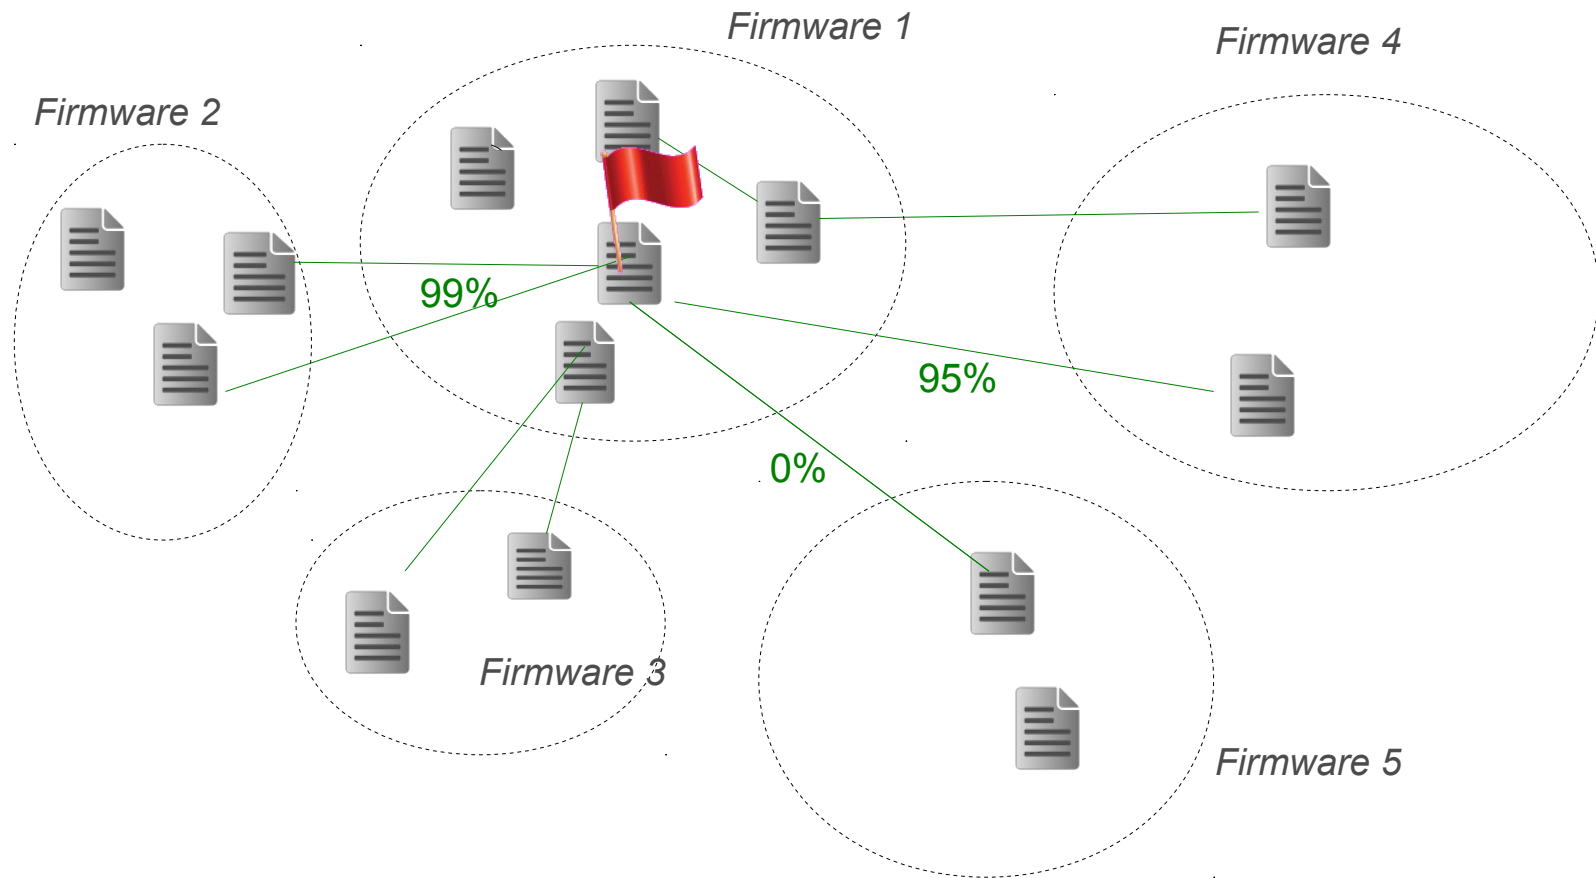
Example: Correlation

- Correlation via fuzzy-hashes (ssdeep, sdhash)



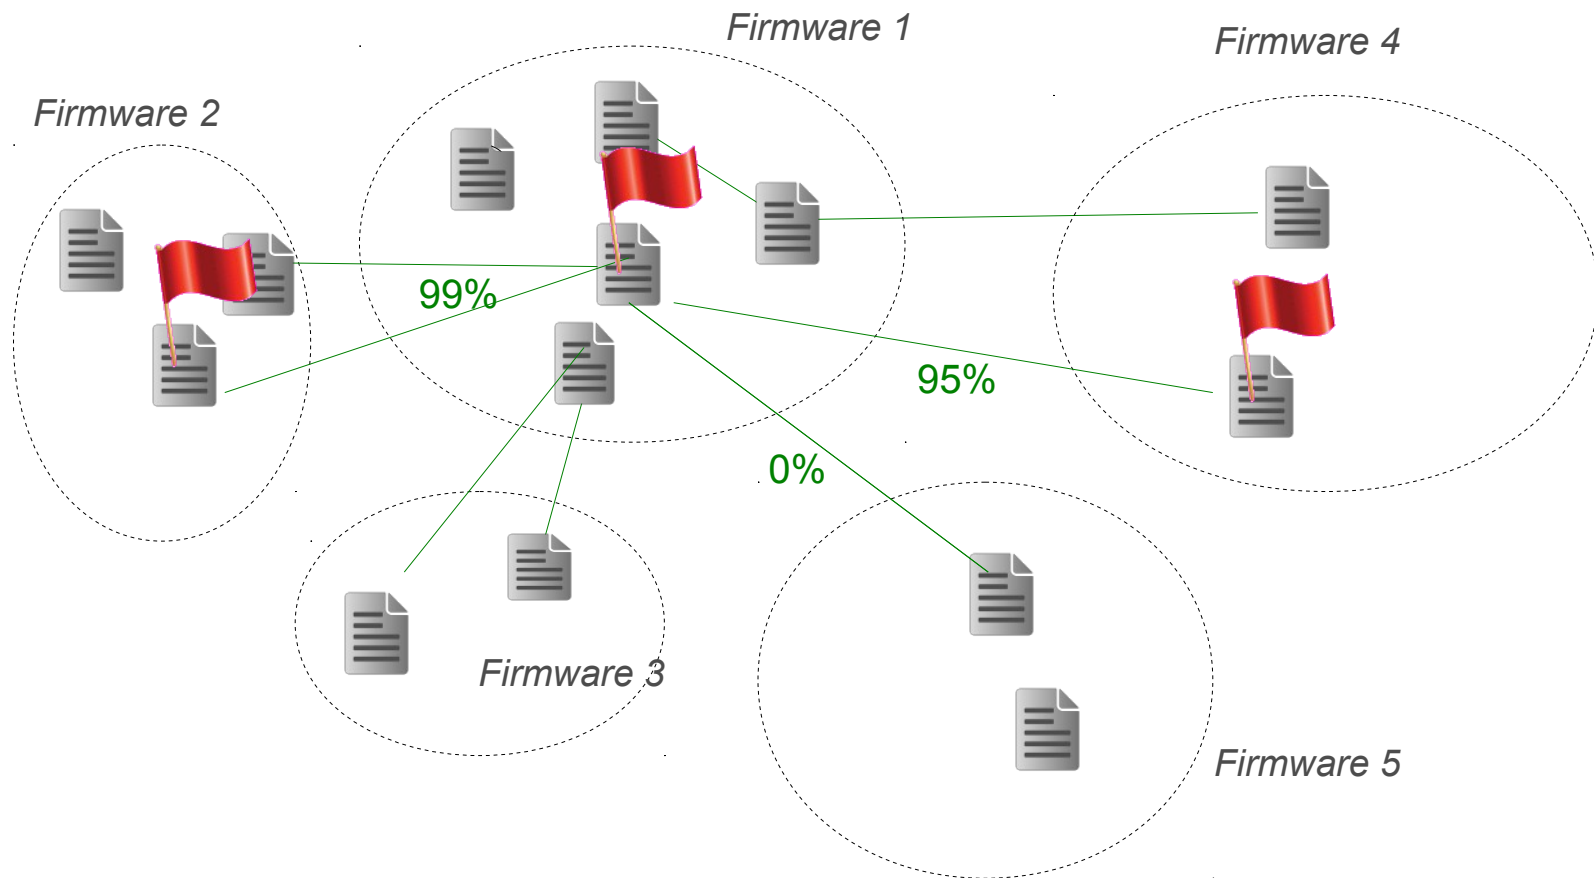
Example: Correlation

- Correlation via fuzzy-hashes (ssdeep, sdhash)



Example: Correlation

- Correlation via fuzzy-hashes (ssdeep, sdhash)




Example: SSL certificates

- SSL cert correlation + vulnerability propagation





Example: SSL certificates

- SSL cert correlation + vulnerability propagation



Example: SSL certificates



- SSL cert correlation + vulnerability propagation



Vendor A

Example: SSL certificates



- SSL cert correlation + vulnerability propagation



Vendor A

Example: SSL certificates

- SSL cert correlation + vulnerability propagation





Vendor A



Example: SSL certificates

- SSL cert correlation + vulnerability propagation





Vendor A




Example: SSL certificates

- SSL cert correlation + vulnerability propagation





Vendor A





Vendor B

Example: SSL certificates

- SSL cert correlation + vulnerability propagation



Vendor A




Vendor B

firmware · əɹ

Example: SSL certificates

- SSL cert correlation + vulnerability propagation



Results: Summary

- 38 new vulnerabilities (CVE)
- Correlated them to 140 K online devices
- Affected 693 firmware files by at least one vuln

”Chamber of Horrors”

- Several recently build images with linux kernels, busybox older than 9 years
- Similar ”debug” backdoor daemon in networking, home automation equipment
- Forgotten or backdoor entries in `authorized_keys` files

”Chamber of Horrors”

- Linux kernel older than 4 years compiled by root on a machine with public IP accepting SSH connections (GPS/Aerospace manufacturer)
- Discovered vulnerability in wireless fireworks system, implemented PoC attack [3]

Contributions Summary

- First large-scale static analysis of firmwares
- Described the main challenges associated
- Shown the advantages of performing a large-scale analysis of firmware images
- Implemented a framework and several efficient static techniques

Conclusions

- A broader view on firmwares
 - Not only beneficial
 - But necessary for discovery and analysis of vulnerabilities
- Correlation reveals firmware relationship
 - Shows how vulnerabilities reappear across different products
 - Could allow seeing how firmwares evolve

Conclusions

- There are plenty of latent vulnerabilities
- Security
 - Tradeoff with cost and time-to-market
 - Clearly not a priority for some vendors

Thank you

- To our advisors, Aurélien and Davide
- To our friends and families
- To the SECURE 2014 organizers
- To everybody who is submitting firmware to us
- To you for listening to this talk :)

firmware • əɪ

The End

Questions?

{name.surname}@eurecom.fr

References

- [1] A. Costin, J. Zaddach, A. Francillon, D. Balzarotti, "A Large-Scale Analysis of the Security of Embedded Firmwares", In Proceedings of the 23rd USENIX Conference on Security (to appear)
- [2] A. Costin, J. Zaddach, "Poster: Firmware.RE: Firmware Unpacking and Analysis as a Service", In Proceedings of the ACM Conference on Security and Privacy in Wireless Mobile Networks (WiSec) '14
- [3] A. Costin, A. Francillon, "Short paper: A Dangerous 'Pyrotechnic Composition': Fireworks, Embedded Wireless and Insecurity-by-Design", In Proceedings of the ACM Conference on Security and Privacy in Wireless Mobile Networks (WiSec) '14