PANDAcap

A Framework for Streamlining Collection of Full-System Traces

Manolis Stamatogiannakis, Herbert Bos, and Paul Groth†

INDE lab†
INtelligent Data Engineering

In this Talk

- Motivation for this work
- Overview of PANDAcap
- Case study: SSH honeypot and dataset
- Conclusion

Motivation

Full-system trace recording

- Log **all instructions** executed and **all data** used.
- Access to full system state – deep analysis.
- Decouples analysis from timing constraints.
- Analysis flexibility.

- Time consuming to setup.
- Very few full-system recording datasets available.

We aspire to lower the barrier for creating full-system recording datasets.

PANDA

- Full System Record + Replay
- Based on QEMU
- Self-contained execution traces
- Analyses implemented as plugins

https://github.com/panda-re

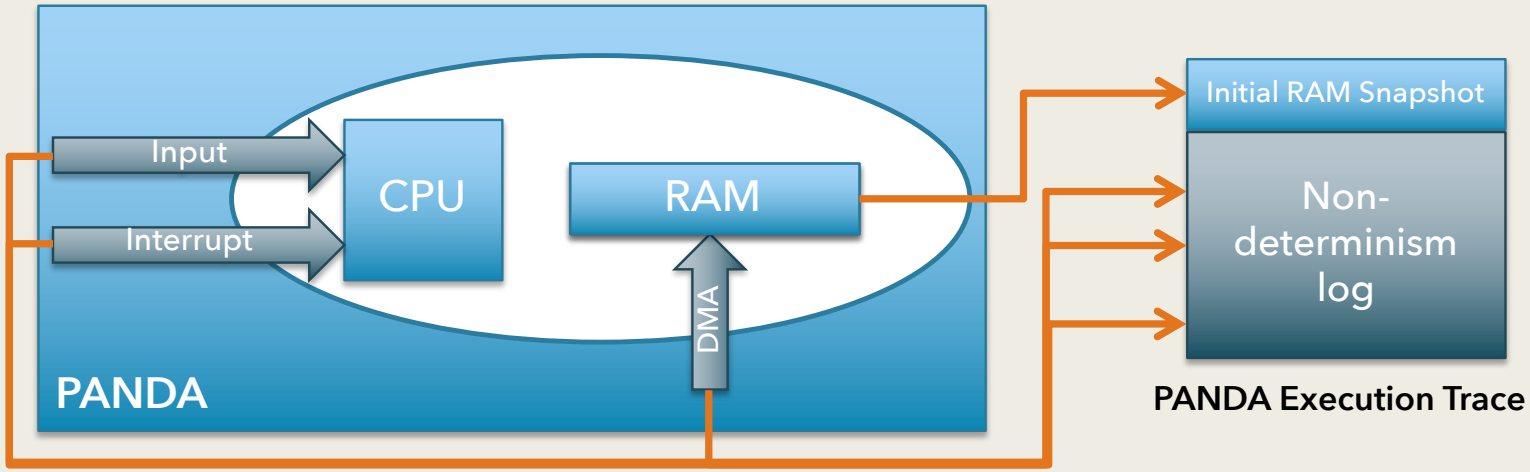

(My) typical PANDA workflow

Prepare for recording

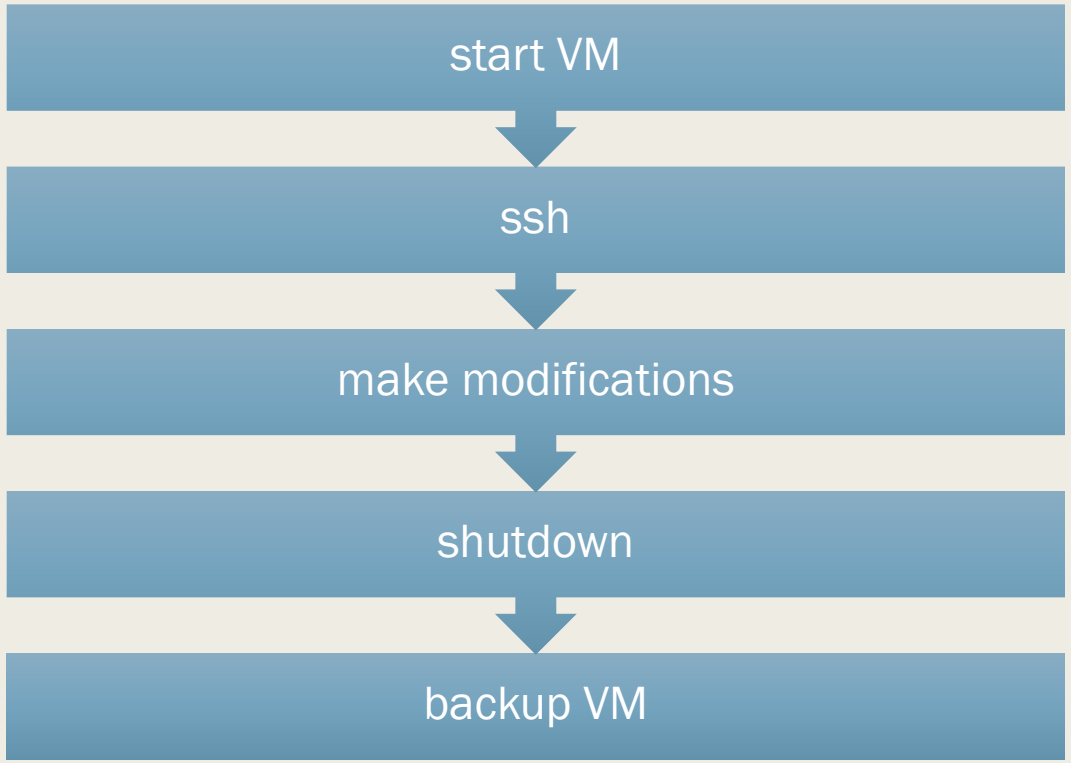

Recording

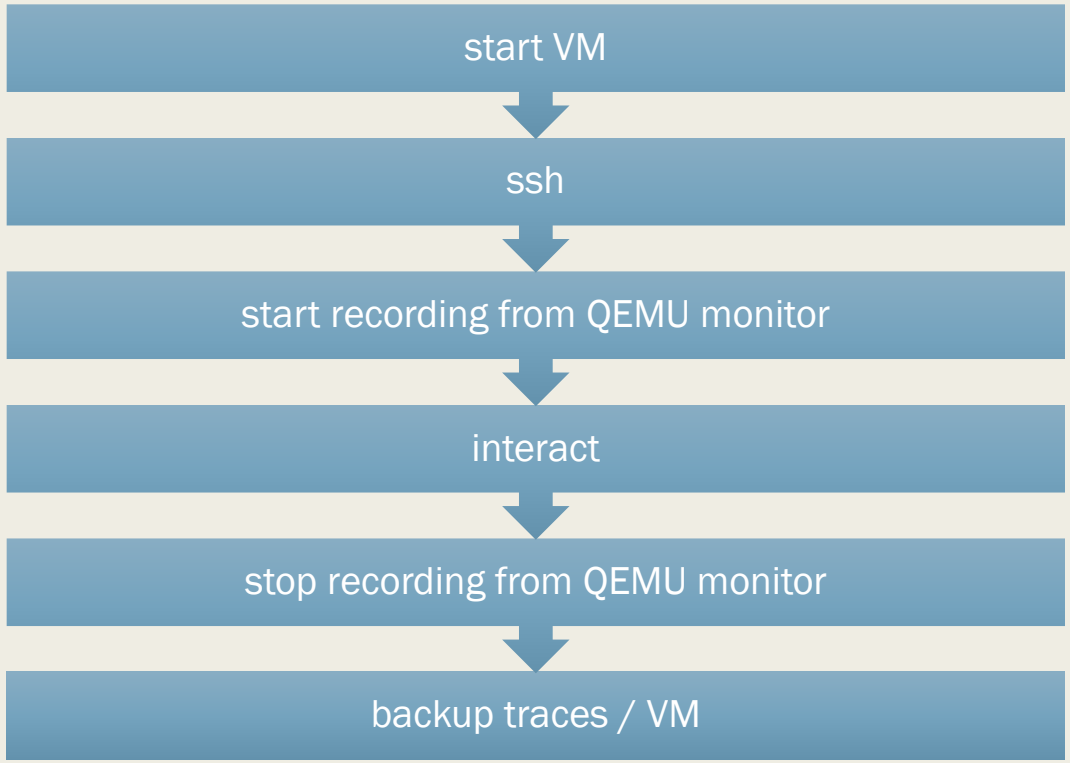

Let's create a PANDA dataset

- The regular PANDA workflow won't cut it.
  - a lot of manual steps
  - error prone (due to the human factor)
- We need to automate things!

Workflow Automation Bottlenecks

- How can I start recording non-interactively?
  - Learn to work with QEMU Monitor Protocol.
- How can I start/stop recording at the right moment?
  - No elegant solution. Bummer!
- How do I move data in/out of the PANDA VM?
  - Deploy ssh keys + sftp?
- How do I replicate the same experiment with different inputs x100?
  - DIY scripting.
- How can I fully utilize my 12 core CPU?
  - …and more DIY scripting.

Now let's put everything together

- Complicated!
- What was it again that I was doing?
- What do you mean I have to start over because I missed X?

MalRec (DIMVA 2018)

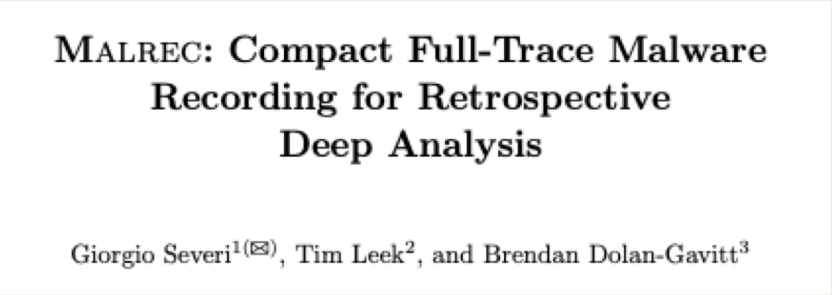

MALREC: Compact Full-Trace Malware Recording for Retrospective Deep Analysis

Giorgio Severi[1(✉)], Tim Leek[2], and Brendan Dolan-Gavitt[3]

- Similar goal with us: create PANDA trace datasets
- Similar approach: off-the-shelf tools
- Purpose-built – not designed to be reusable.

  "This is not intended to work for anyone else out of the box, just to provide a starting point. You will undoubtedly have to make heavy local modifications."

- Last update in 2015 – tooling hasn't been modernized since.

Fast forward to 2020

- Containers are mainstream.
  - networking virtualization
  - storage virtualization
  - ease of deployment
- Some containers available for PANDA
  - geared towards testing builds
- Runtime customization of PANDA VMs still a DIY affair.

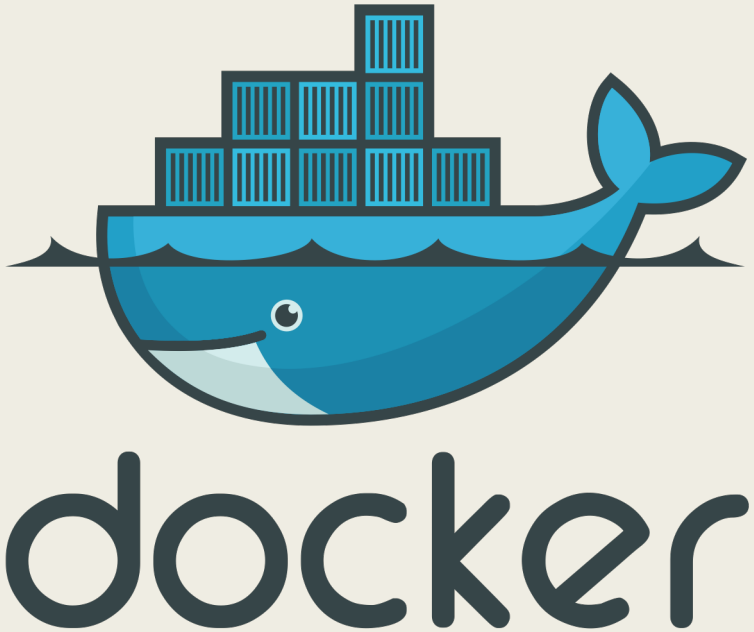

We can improve on this.

PANDAcap Overview

Enter PANDAcap

- Accurate start/stop of recording.
- Supports Docker – lean image.
- Streamlined VM bootstrapping.
  - rc.d-like initialization process
  - Jinja2 templating support
- Command line wrapper providing access to most commonly used features of Docker/PANDA.

The recctrl plugin

- Accurate start/stop of recording.
- Building block: PANDA_CB_GUEST_HYPERCALL.
- Support for sessions (semaphore-like).
- Support to specify the PANDA recording name from the guest.
- A timeout can be specified for limiting the length of the recording.
- Batteries included: recctrlu guest utility

Lean Docker Image

- Contains only runtime dependencies.
- Bootstrapping mechanism for Docker runtime environment.
- Shared configuration with VM runtime bootstrapping.
- Mountpoints affecting a run:
  - Docker runtime bootstrap directory
  - QCOW image for PANDA
  - Recording output directory
  - X11 server path

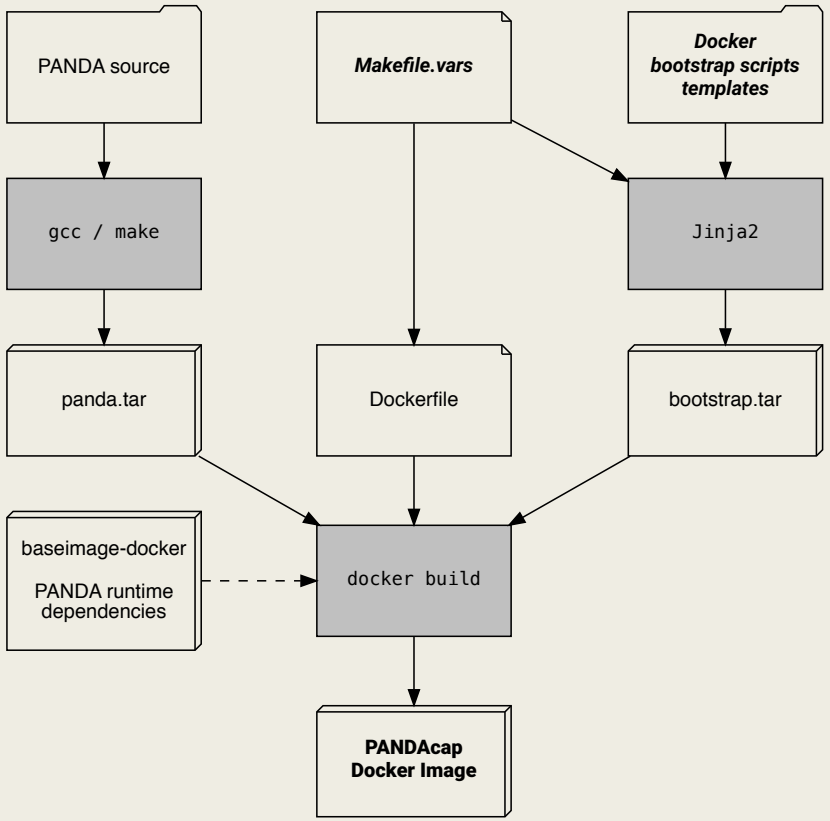

Runtime bootstrapping – layout

bootstrapping scripts

files used by the scripts

environment template / Makefile

Makefile targets

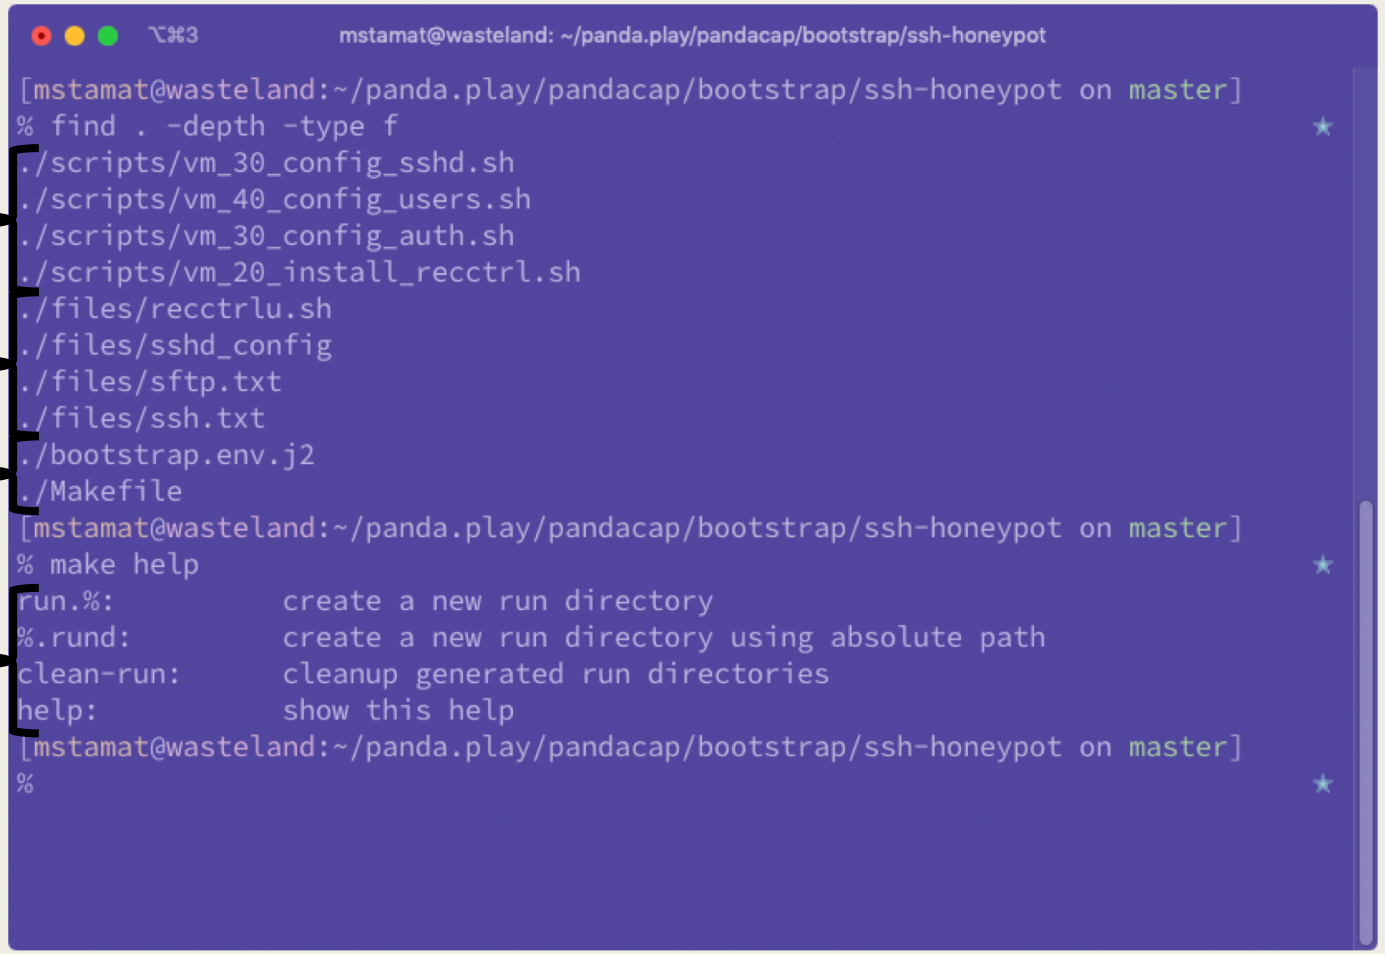

Runtime bootstrapping – output

VM runtime bootstrapping

Docker runtime bootstrapping

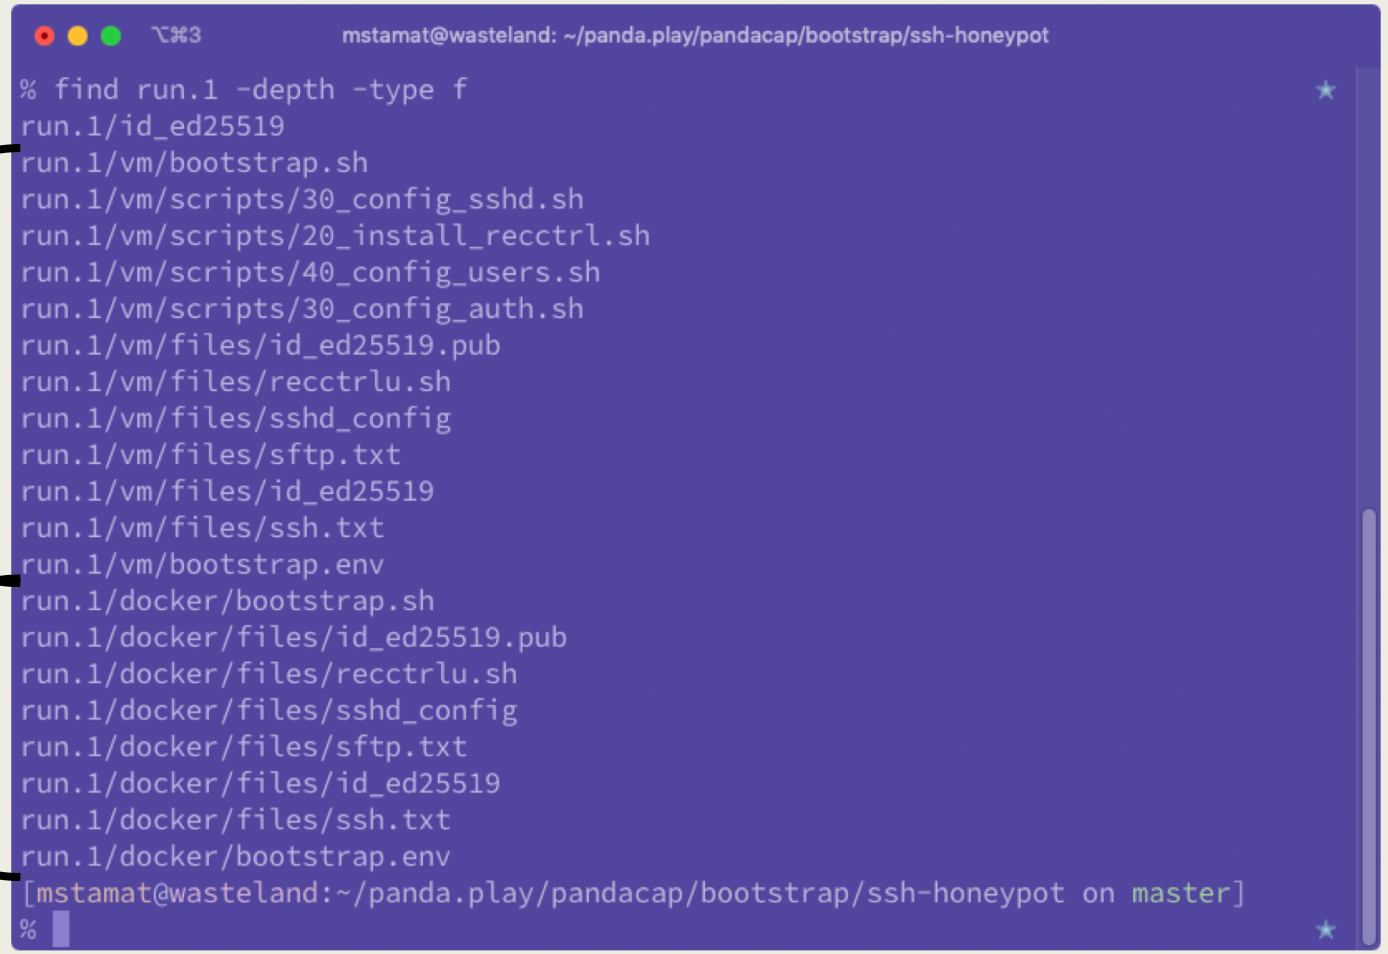

pandacap.py wrapper

```
⌥⌘2                    mstamat@wasteland: ~
[mstamat@wasteland:~]
% /opt/panda/bin/panda-system-i386 --help | grep '^ *-' | wc -l
179
[mstamat@wasteland:~]
% docker run --help | grep '^ *-' | wc -l
93
[mstamat@wasteland:~]
%
```

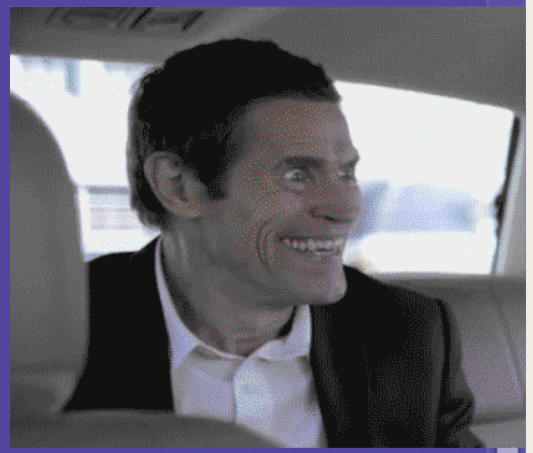

Most common PANDA/Docker options

PANDA

- Disk configuration.
- Network configuration and port forwarding.
- Creation of delta image.*
- Creation of bootstrap disk.*
- Memory/Arch configuration.
- Display configuration.

* Involves additional tools.

Docker

- Mount configuration.
- Network configuration and port forwarding.

pandacap.py wrapper

pandacap.py wrapper

- All common options in one place.
- Takes care of:
  - Creation of bootstrap disk for the VM.
  - Initialization of a new delta image for the VM.
  - Proper escaping of commands.
- Output files/images are labeled so concurrent runs can be told apart.
- Does not mandate the use of Docker.
  - Can be used as a simple wrapper around PANDA.

PANDAcap source code

github.com/vusec/pandacap

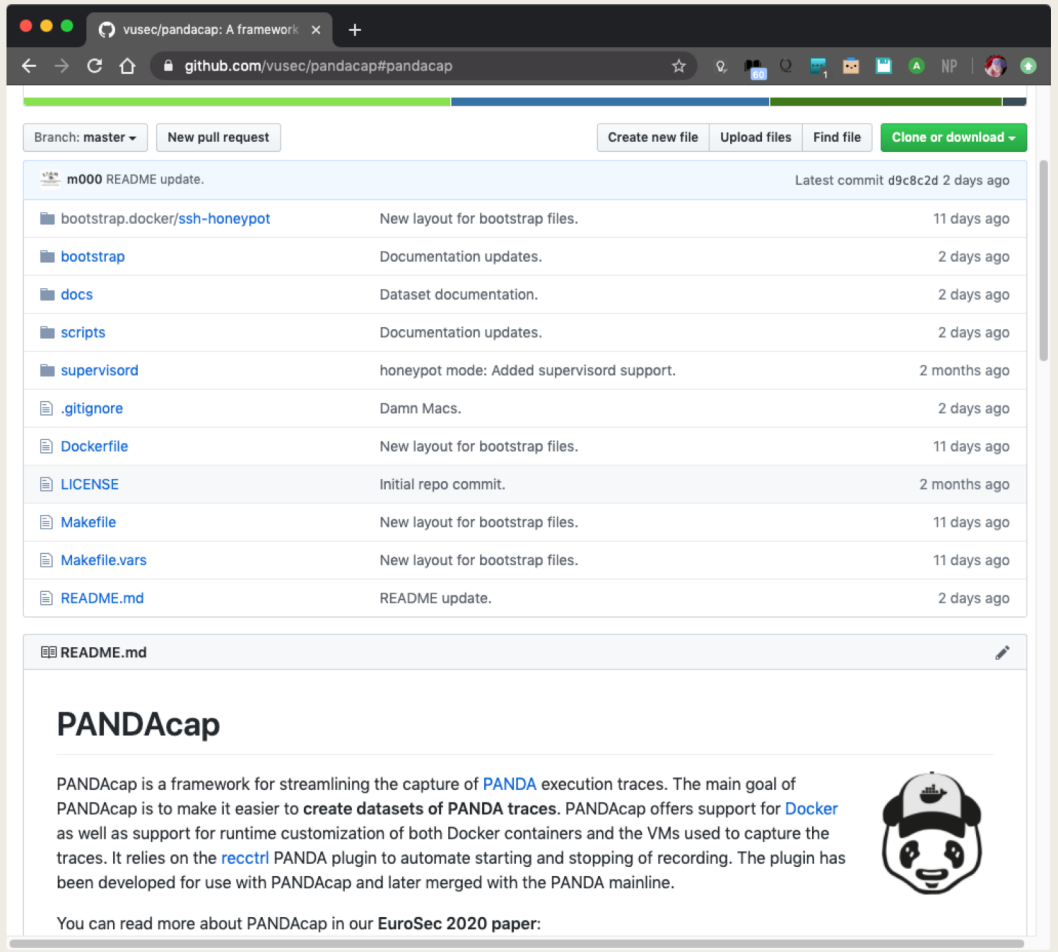

Case Study: SSH Honeypot

and dataset

PANDAcap Case Study: ssh honeypot

- Brute-force ssh attacks are still popular.
- In their 2016 survey of existing honeypot software, Nawrocki et al. mention no honeypot based on full system Record and Replay. https://arxiv.org/abs/1608.06249
- Full system Record and Replay offers significant advantages:
  - Flexibility of analysis.
  - Captures all transient effects on the system.

- Common misconception: Analyzing an ssh intrusion is trivial.

In a Slack channel somewhere…

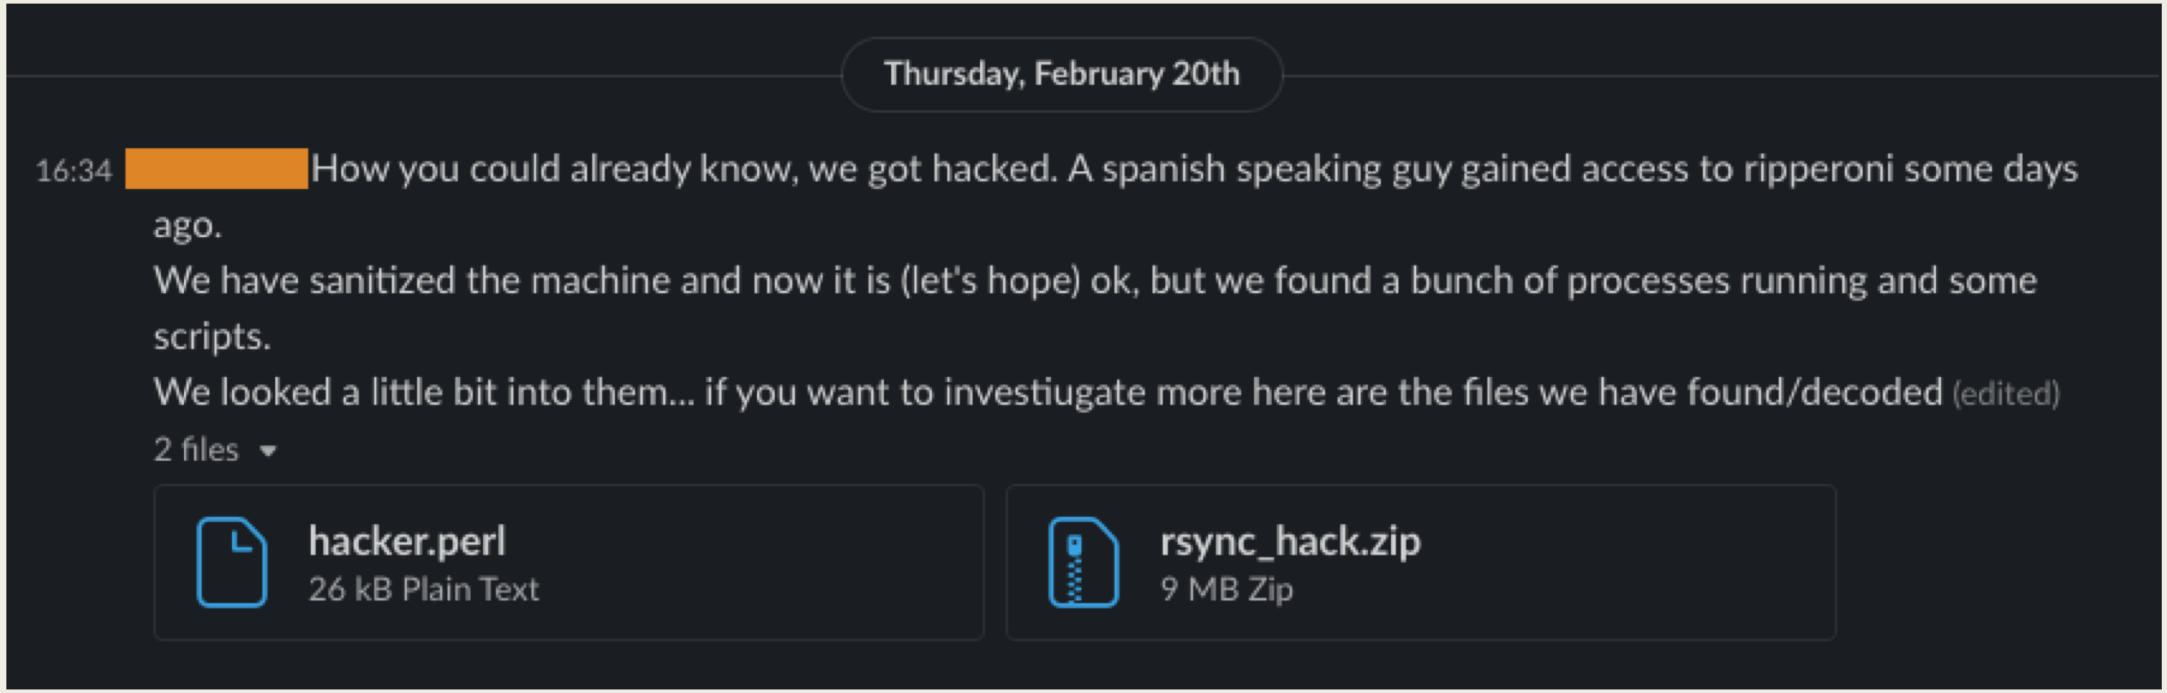

In a Slack channel somewhere…

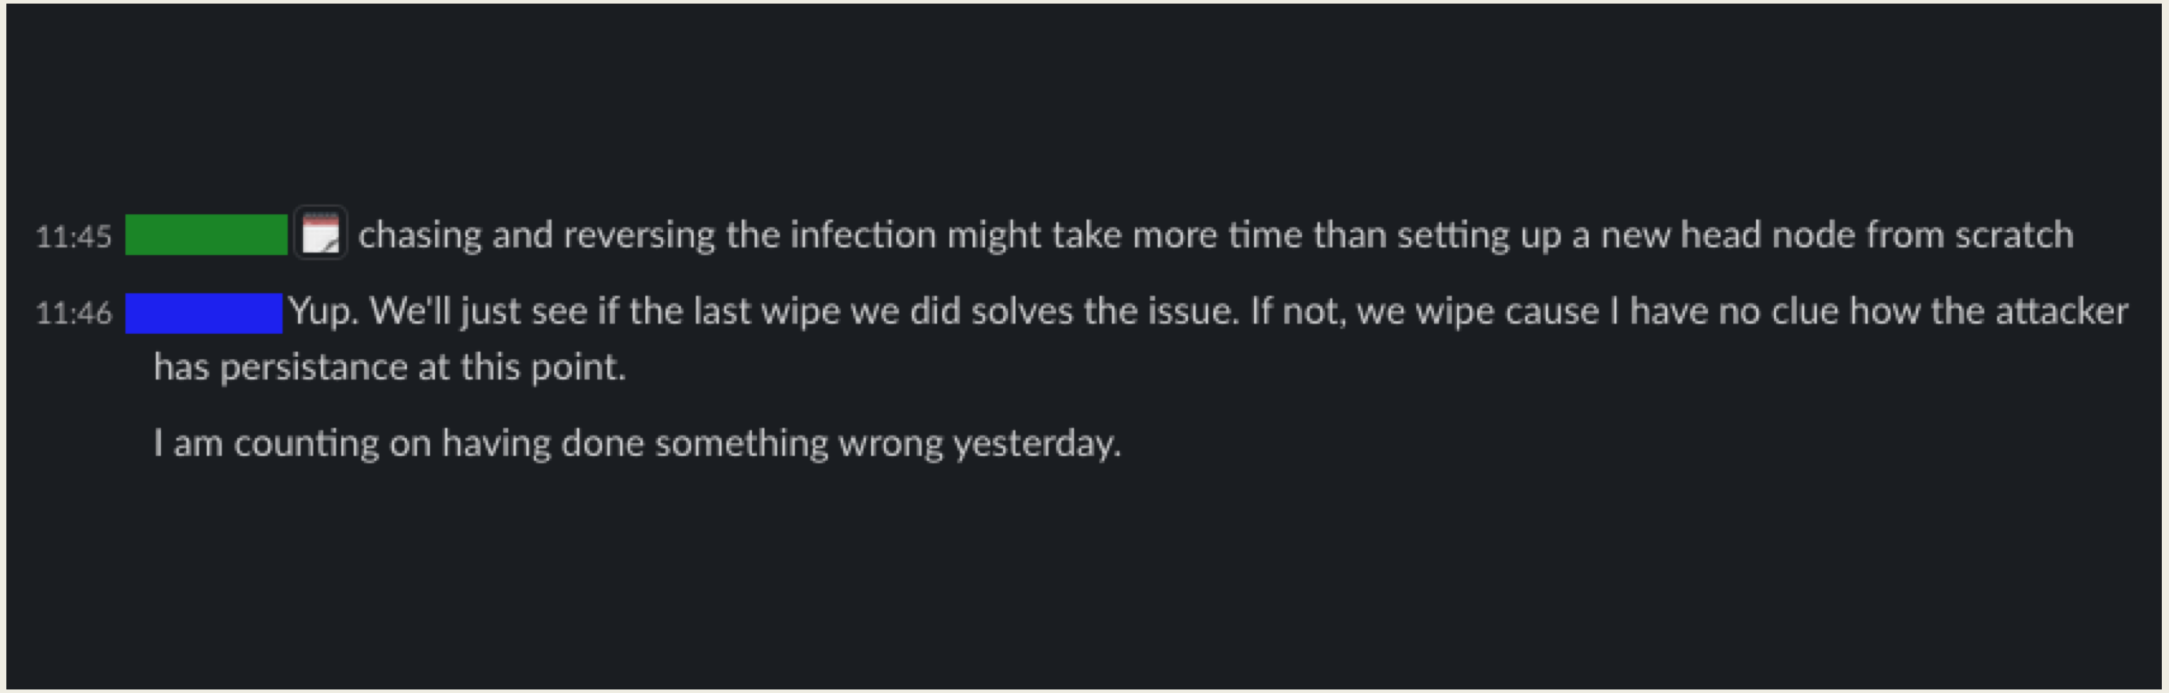

In a Slack channel somewhere…

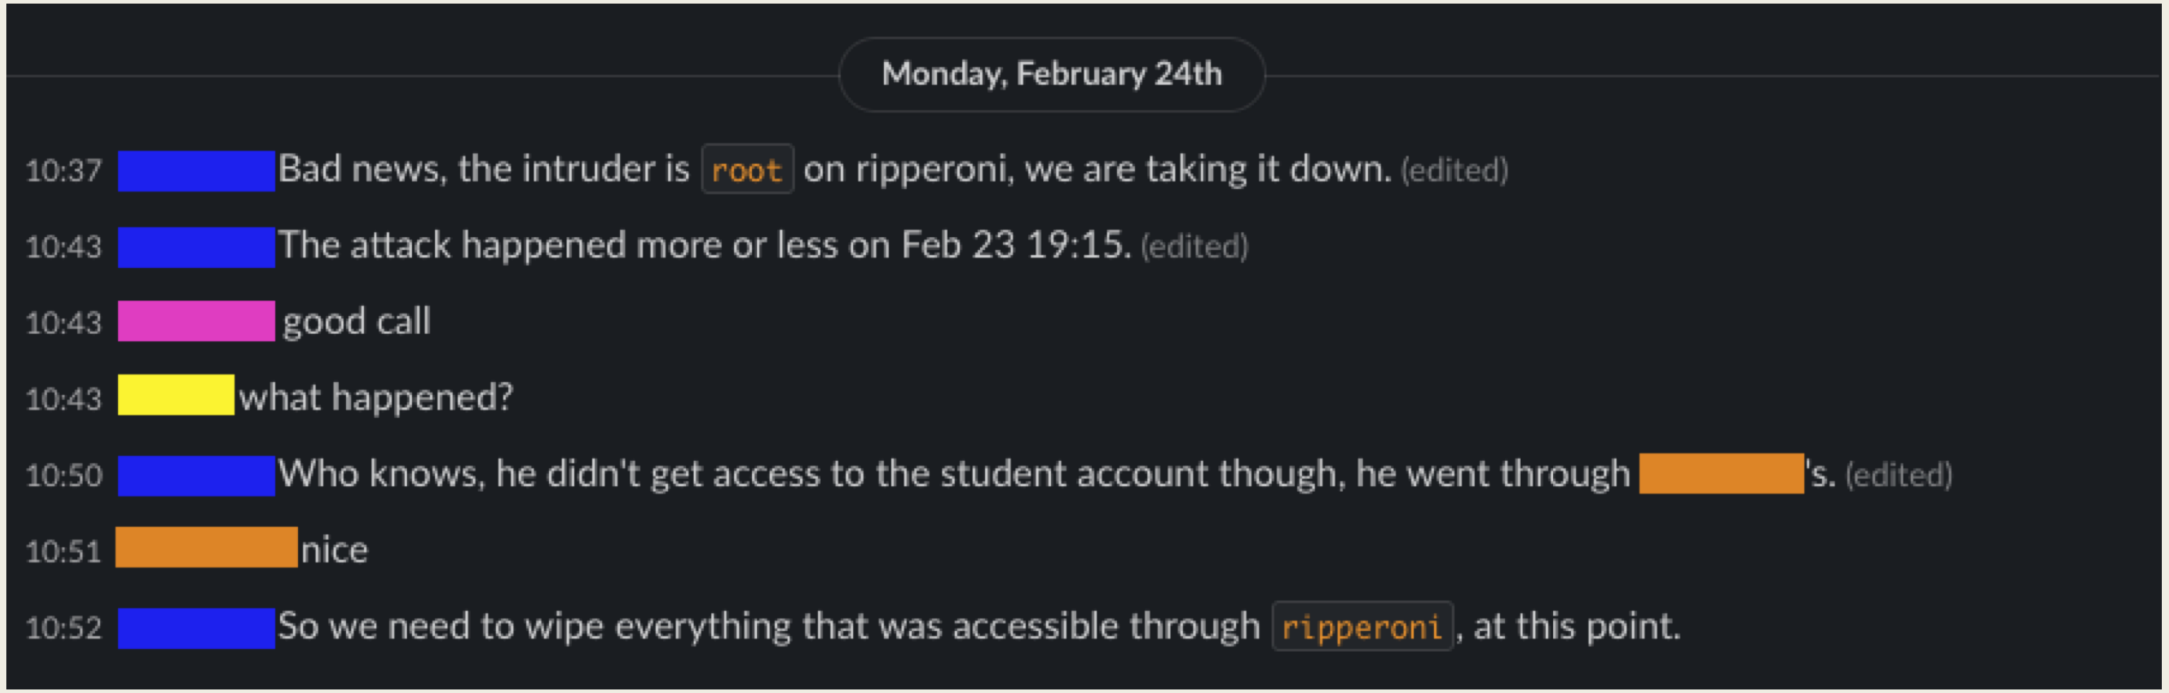

Aftermath

- No point of entry was determined.
- Unsure how privilege escalation was achieved.
- Partial recovery of the hacker's tools.
- Partial log of communications.
- Failed to cleanup the machine properly.

- **Post-mortem analysis is hard, even for experts.**
- PANDA system-tracing can provide answers!

Honeypot analysis with PANDA

- Privilege escalation → exact trace of system calls that led e.g. to a privileged execve
- Hacker tools → ability to fully reconstruct them from the non-determinism log, even if they have been “shredded”
- Communication logs → pcap files + access to unencrypted network stack buffers
- Cleaning up the system → produce a detailed provenance log for all the files that were modified, identify potentially malicious modifications

PANDAcap honeypot dataset

- Ran the experiment for ~3 days on a single IP address.
- Traces limited to 30′.
- Out of 3 ports used, only 2 were visited.
- Collected 63 traces in total.
- Compressed size (including disk deltas) ~23Gb.

Table 1: Collected samples per *ssh* port. No attempts to gain access to the VM listening on port 2200 were made.

port	samples	nondet	nondet-gz	disk-delta
22	50	9.61 GiB	2.75 GiB	11.49 GiB
2222	13	0.99 GiB	0.28 GiB	3.00 GiB

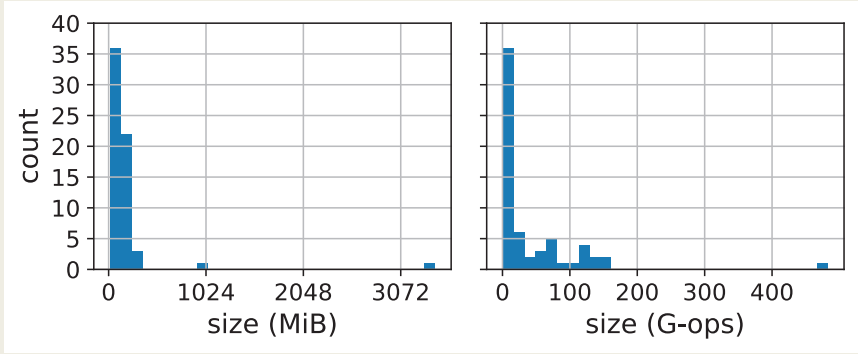

Figure 2: Trace size and instruction count distributions.

PANDAcap honeypot dataset

- Quick qualitative analysis revealed a variance of behaviours.
- Different roles:
  - SSH scanning vs. HTTP/S communication
- Different “return” patterns:
  - 2 logins was the most common case
  - 68 logins was the most common
  - only 2 instances of full log wiping

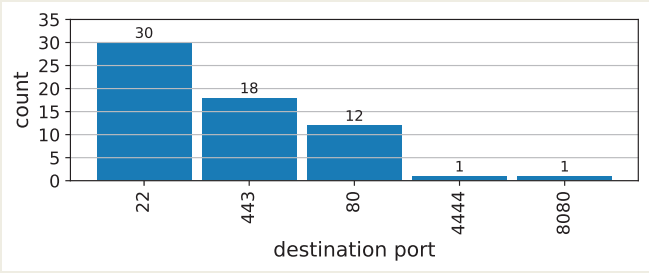

Figure 3: Top target ports for outgoing connections. In one trace, there were no outgoing connections.

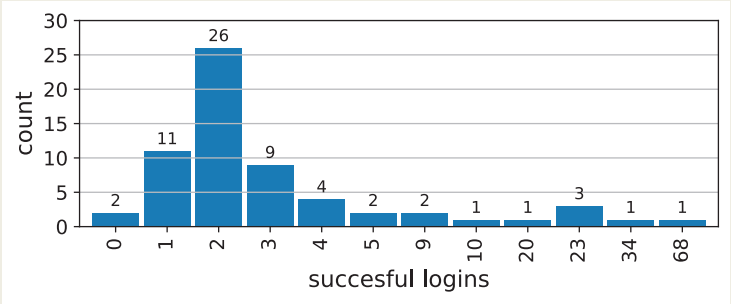

Figure 4: Succesful logins attempts in auth.log.

PANDAcap honeypot dataset availability

zenodo.org (CERN)

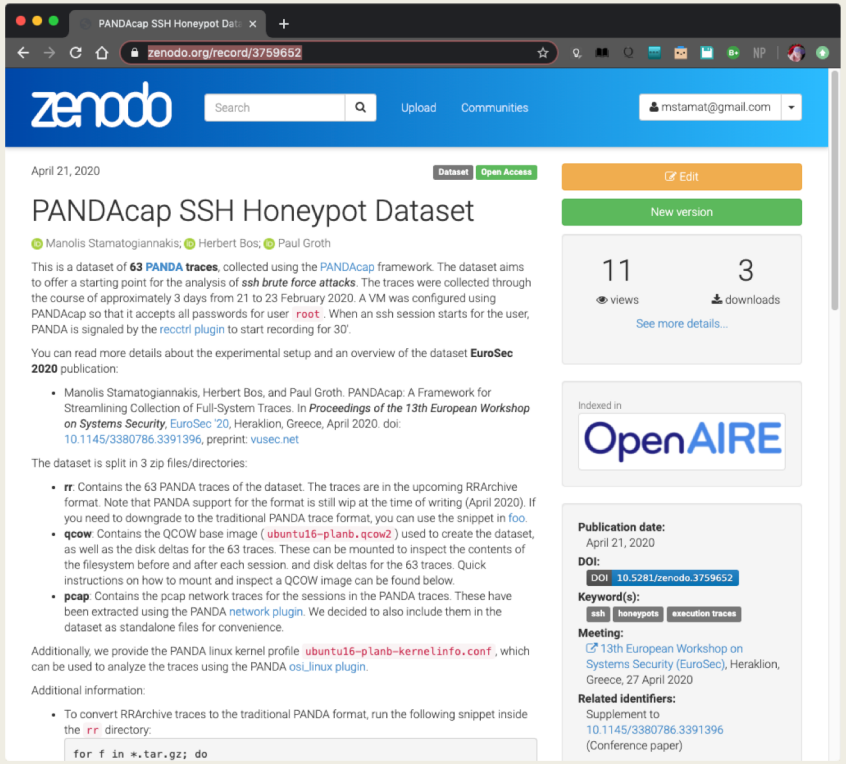

academictorrents.com

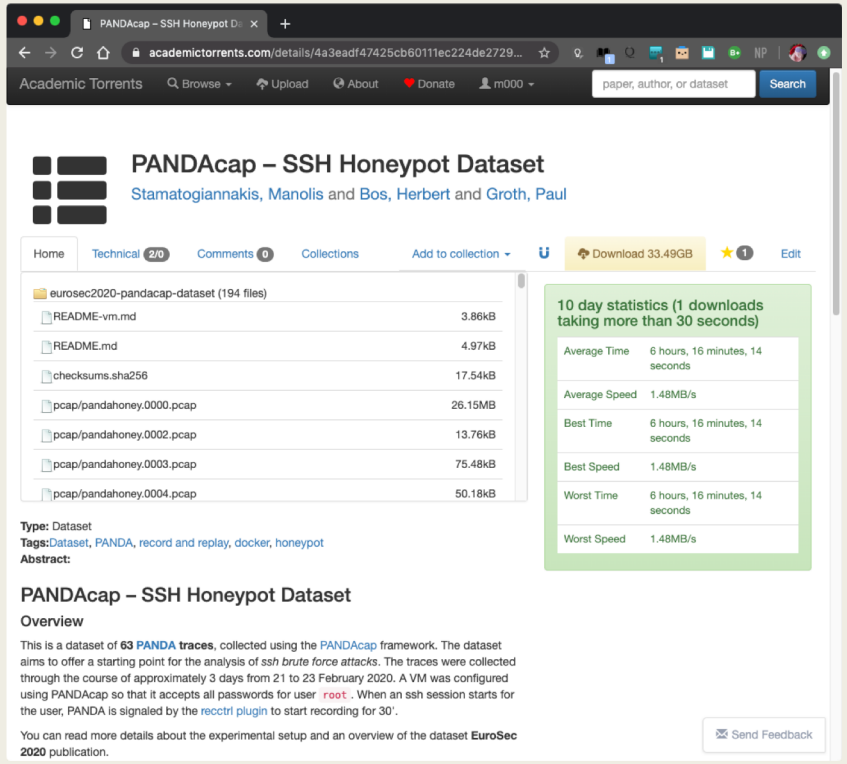

Conclusion

Conclusion

- PANDAcap:
  - easier creation of PANDA trace datasets
  - Docker support
  - streamlined bootstrapping
  - Apache 2.0 license
- PANDAcap SSH honeypot dataset:
  - 63 samples
  - CC 4.0 license

More Information

Code & dataset

Twitter

#PANDAcap #eurosec2020

@vusec

@inde_lab_ams

github.com/vusec/pandacap