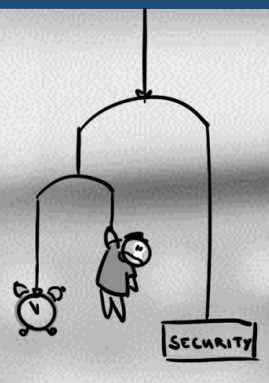

Reflections on Active Network Measurements in Academia

(And the others, too. A bit.)

Tobias Fiebig
Max-Planck Institut für Informatik

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Network Measurements

- Network Measurements:
  The thing we do
- Important tool for academics (getting papers) and practitioners (getting something useful to improve protocols)
- Come in active or passive
- Especially active ones are difficult to do well

© Constanze Dietrich

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

supported by:
RIPE NCC Community Projects Fund
LWLCOM connecting you.
OPENFACTORY
Virtuacloud
DE-CIX
WIIT THE PREMIUM CLOUD

Academia

- The ‘place’ where we go to study
- Academics: The people doing the studying and teaching
- Different levels of education:
  - Bachelor: Show you can do stuff you are told
  - Master: Show you can do stuff without being told what to do
  - PhD: Show that you can come up with stuff and then do it yourself
    (Also: Usually starting after the master, and the ones doing most research)
- Purpose:
  - Find new knowledge & technology
  - Make the world a better place
  - Educate people to go and make the world a better place
- Main currency: Research output, Papers, “Renown & Fame”

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Why I can talk about academia…

- B.Sc. Cognitive Science (2012)
- M.Sc. System & Network Engineering (2013)
- Dr.-Ing. Network Measurement and IT Security (2017)
- Until March 2022 Assistant Professor at TU Delft
- Since April 2022 Max-Planck-Institut für Informatik

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

… and I even know a bit about ops.

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Disclaimer

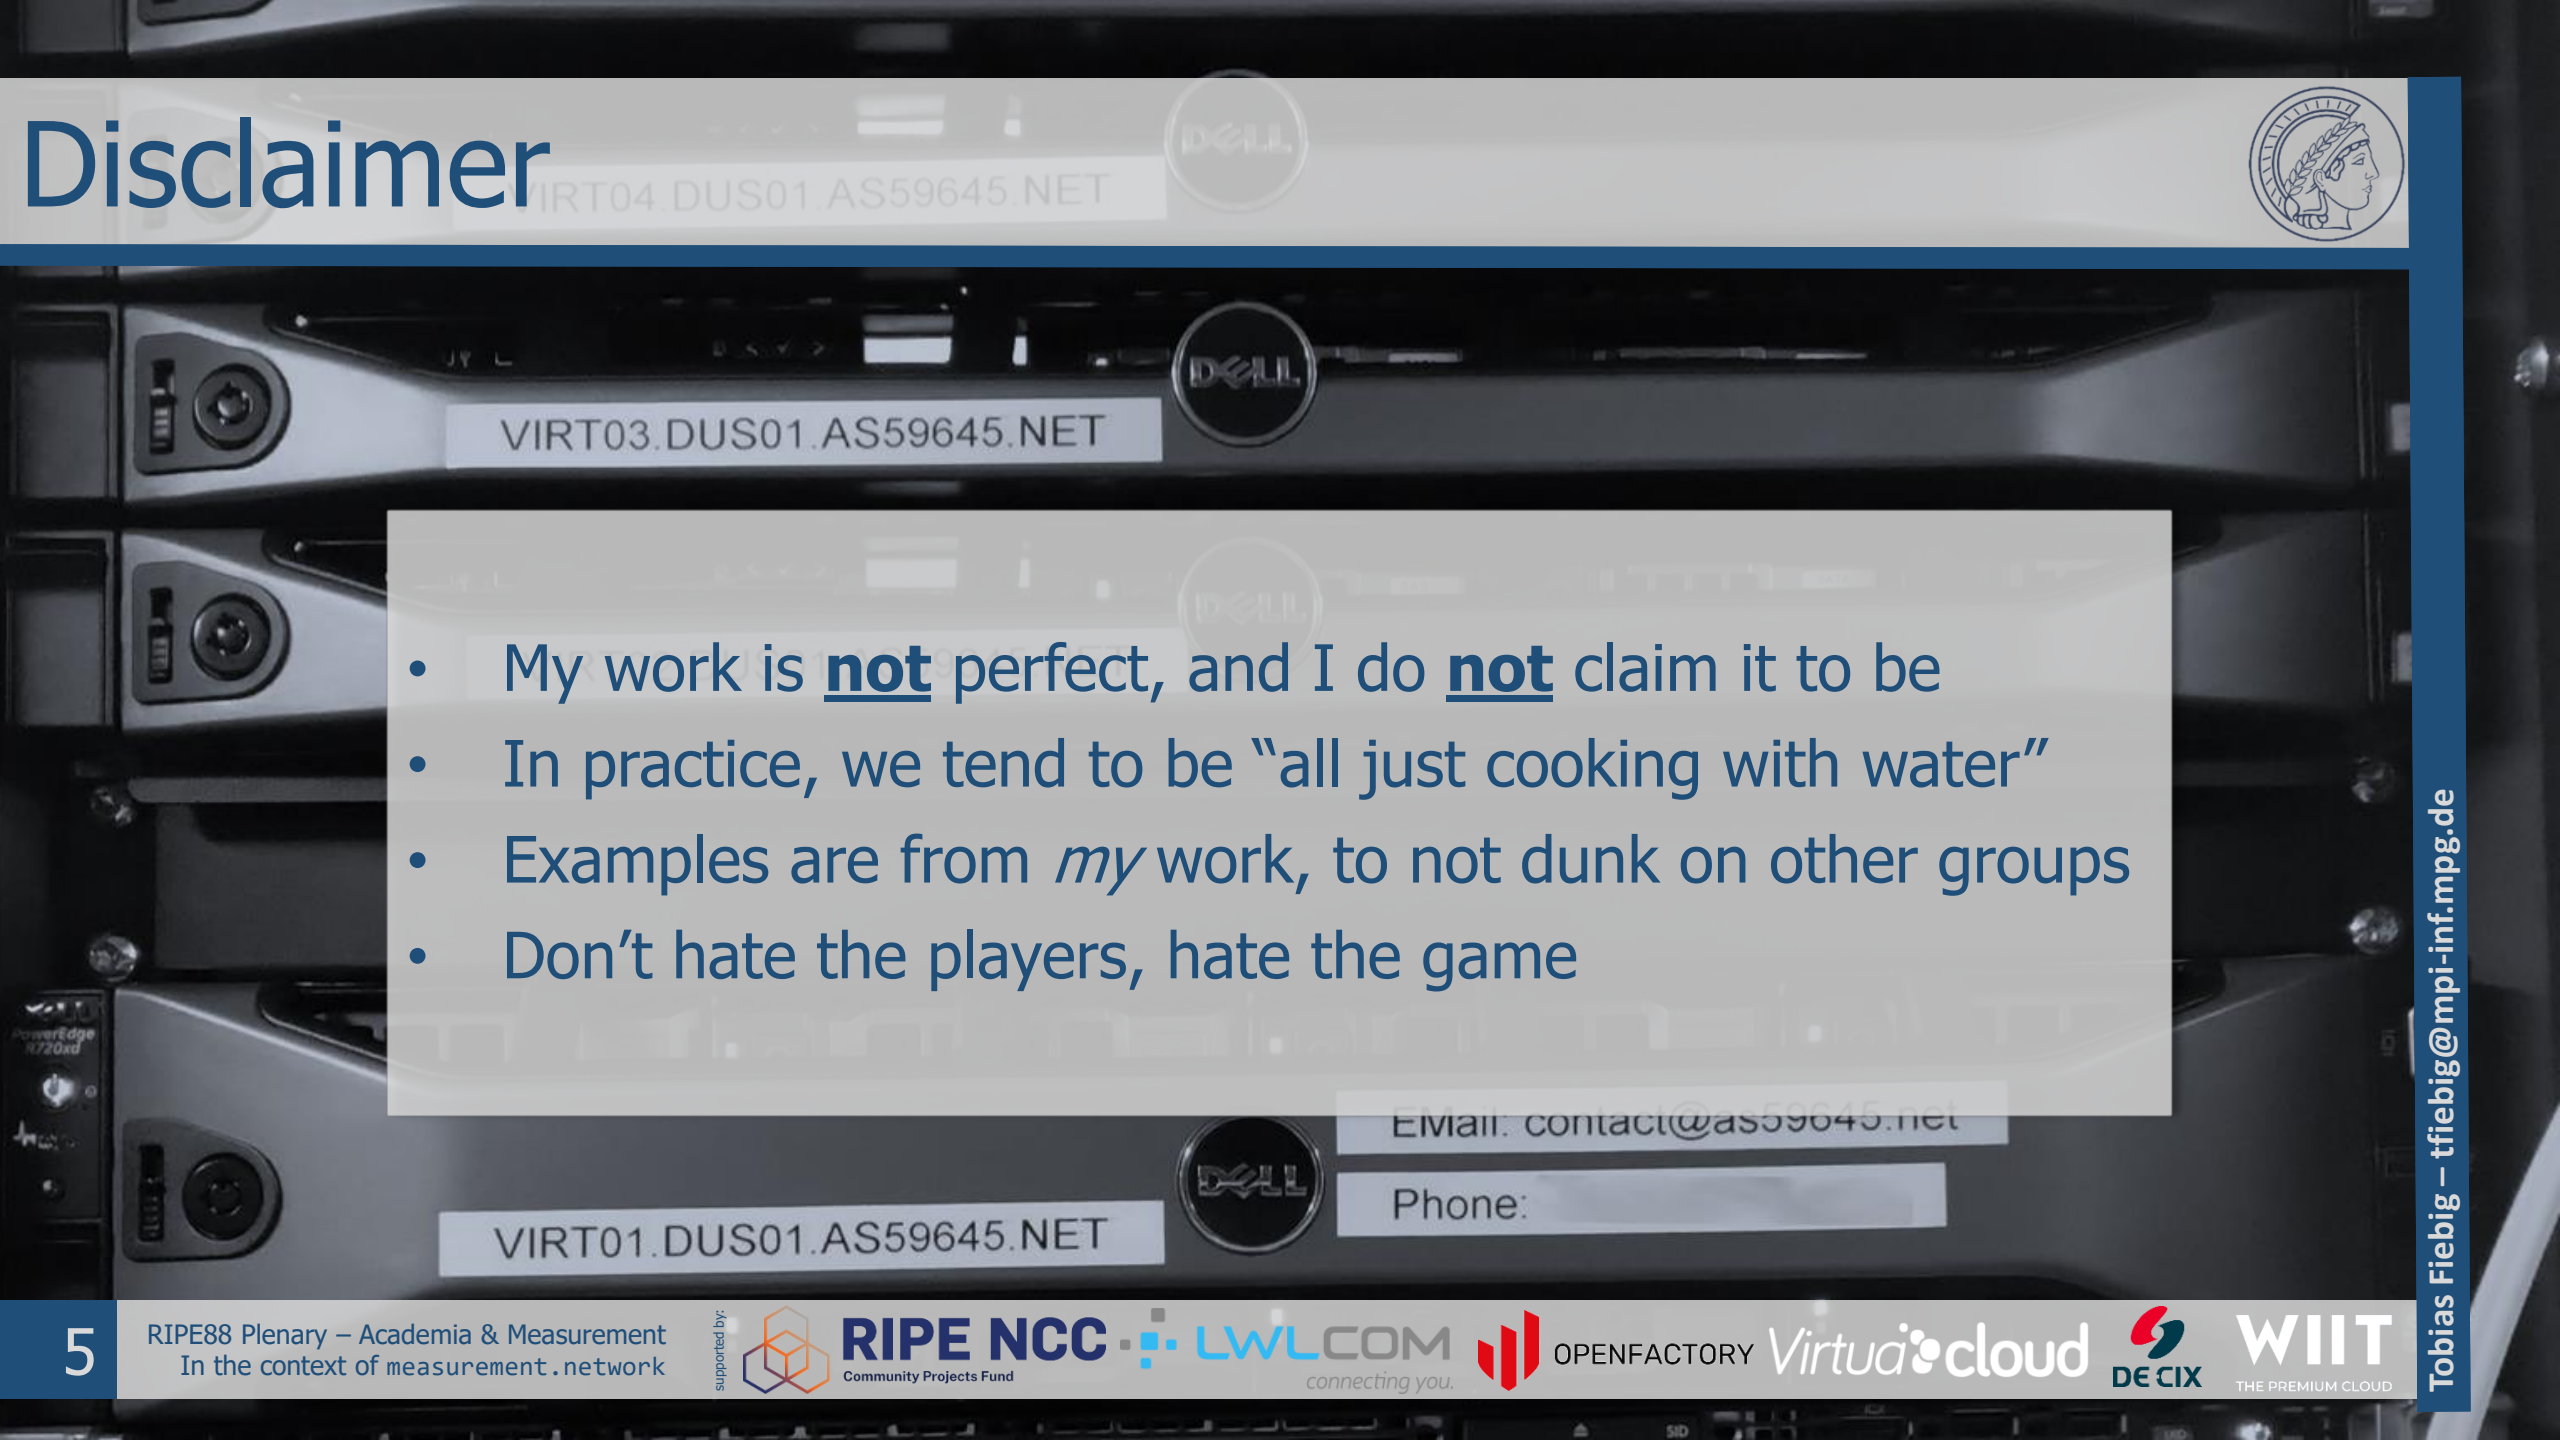

- My work is **not** perfect, and I do **not** claim it to be
- In practice, we tend to be "all just cooking with water"
- Examples are from *my* work, to not dunk on other groups
- Don't hate the players, hate the game

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

(Me) Doing Stupid Things

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

(Me) Doing Stupid Things

Current Status:	CRITICAL (for 0d 0h 0m 7s)
Status Information:	connect to address mail.aperture-labs.org and port 25: Connection refused SMTP CRITICAL - 0.003 sec. response time

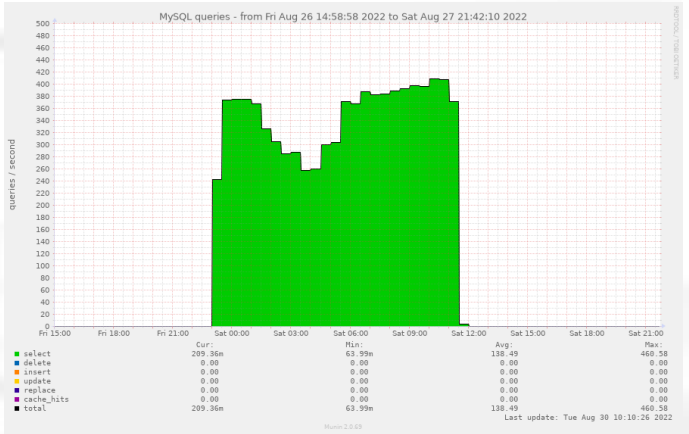

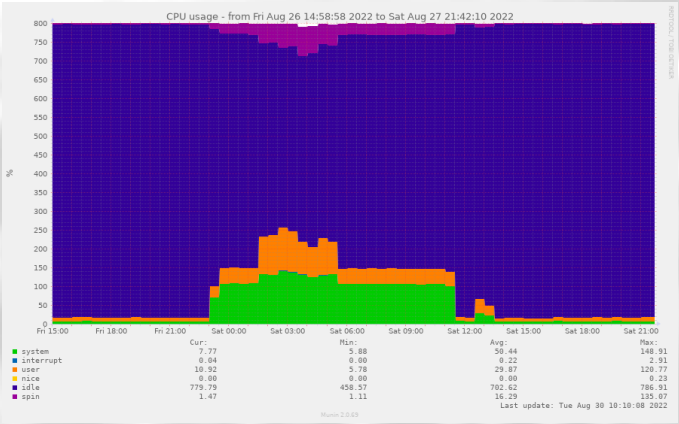

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

(Me) Doing Stupid Things

No SMTP

Current Status:	CRITICAL (for 0d 0h 0m 7s)
Status Information:	connect to address mail.aperture-labs.org and port 25: Connection refused SMTP CRITICAL - 0.003 sec. response time

MySQL @ 400 qps

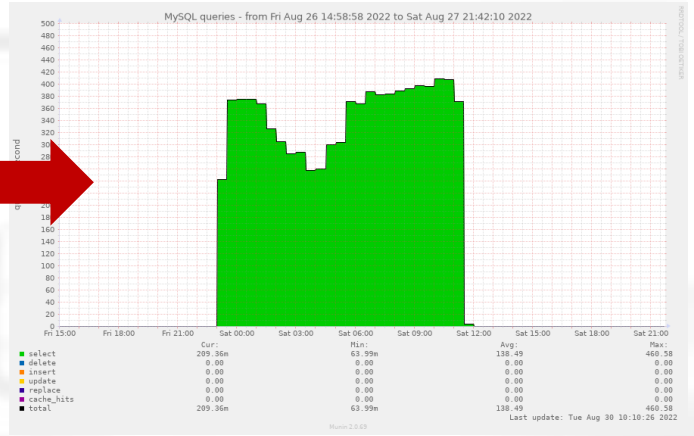

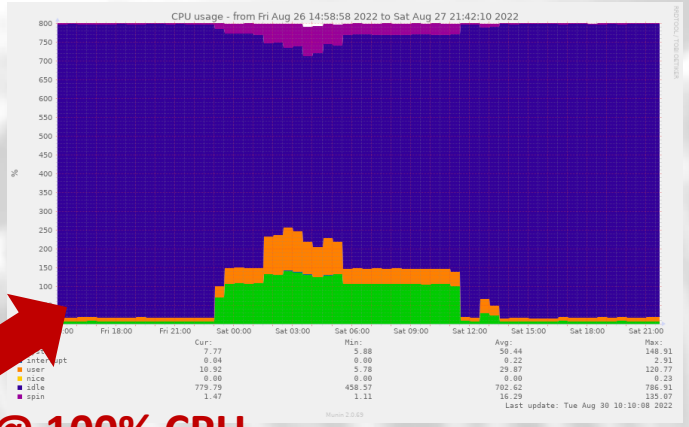

OpenSMTPd @ 100% CPU

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

(Me) Doing Stupid Things

No SMTP

Current Status:	CRITICAL (for 0d 0h 0m 7s)
Status Information:	connect to address mail.aperture-labs.org and port 25: Connection refused SMTP CRITICAL - 0.003 sec. response time

```
          mysql_native_password\0“
91    54126 table-mysql RET   recvfrom 92/0x5c
92    54126 table-mysql CALL  sendto(3,0xf2204a82000,0xcc, \
          0x480<MSG_DONTWAIT|MSG_NOSIGNAL>,0,0)
93    54126 table-mysql GIO   fd 3 wrote 204 bytes
94          "\M-H\0\0\^A\M^L\M-"\M->\0\0\0\0@-\0\0\0\0\0\0\0 \
          \0\0\0\0\0\0\0\0\0\0\0\^]\0\0\0opensmtpd\0 [...] \
          smtpd\0mysql_native_password\01\^C_os\a \
          OpenBSD\f_client_name
95          libmariadb\^D_pid\^E54126\^O_client_version\^E3.3.2 \
          _platform\^Eamd64\f_server_host 127.0.0.1“
96    54126 table-mysql RET   sendto 204/0xcc
...
147         "t\0\0\^A\M^?\M-s\^D#HY000Illegal mix of collations \
          (latin1_swedish_ci,IMPLICIT) and (utf8mb4_general_ci, \
          COERCIBLE) for operation '=“
...
153         "credentials[54126]: warn: trying to reconnect after \
          error: Illegal mix of collations (latin1_swedish_ci, \
          IMPLICIT) and (utf8mb4_general_ci,COERCIBLE) for \
          operation '='
```

MySQL @ 400 qps

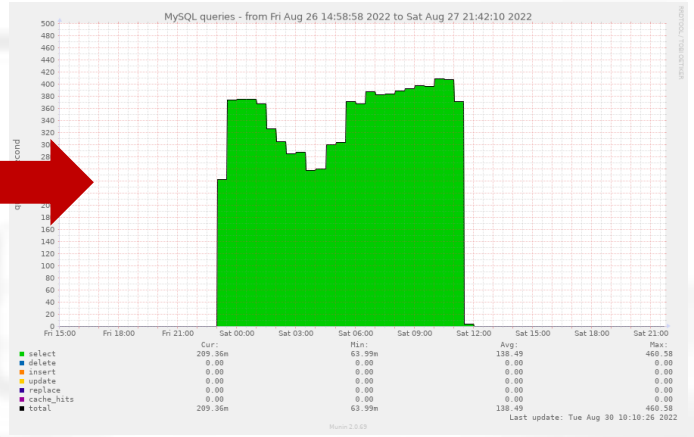

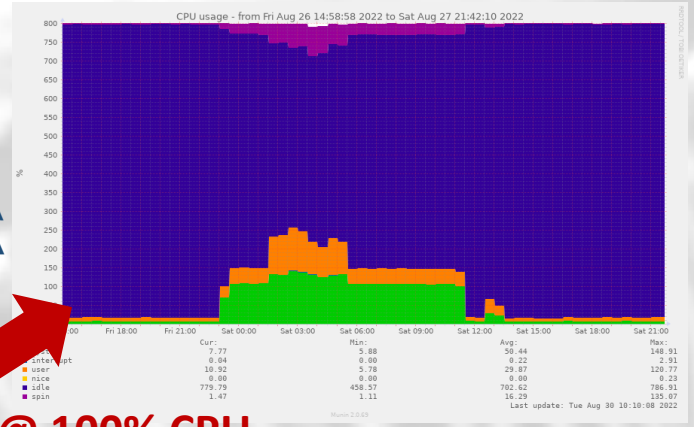

OpenSMTPd @ 100% CPU

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Measuring Email Security

- There is 'SPF' (Sender Policy Framework):
  - "`v=spf1 ip4:192.0.2.1 include:spf.example.com -all`"
- 'include' smells like 'recursion' and 'recursion' smells like 'DoS'
  - How much time would mailservers spend traversing an eternal SPF tree? The RFC says 'stop after 10', so for sure EVERYONE does, right?
- You smell a paper, get excited, and start some measurements
- Quick Hack: Python DNS server & go!

- => Some mailservers actually *do* keep traversing a tree of eternal depth…

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

supported by:
RIPE NCC Community Projects Fund
LWLCOM connecting you.
OPENFACTORY
Virtuacloud
DE-CIX
WIIT THE PREMIUM CLOUD

Making People Unhappy

- You recently moved institutions and the reverse/forward DNS settings for probe identification / website fell victim to it
- People couldn't reach you

- You approach mail operators about your work and they are "not necessarily really happy" with what you did…

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

supported by: RIPE NCC Community Projects Fund

Running Reliable Network Measurements

```
root@msrmnt.example.com:/opt/yolo-colo# docker compose up -d
root@msrmnt.example.com:/opt/yolo-colo# docker ps
CONTAINER ID   IMAGE            COMMAND                  CREATED         STATUS         PORTS      NAMES
33ecc4cc2bfb   yolo-colo/scan   "/docker-entrypoint.…"   21 hours ago    Up 21 hours    25/tcp     scan_1
8b86f724e188   yolo-colo/rdns   "/bin/sh -c /entrypo…"   21 hours ago    Up 21 hours    53/udp     rdns_1
15d771655355   yolo-colo/rdns   "/bin/sh -c /entrypo…"   21 hours ago    Up 21 hours    53/udp     adns_1
```

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Running Reliable Network Measurements

```
root@msrmnt.example.com:/opt/yolo-colo# docker compose up -d
root@msrmnt.example.com:/opt/yolo-colo# docker ps
CONTAINER ID   IMAGE             COMMAND                  CREATED         STATUS         PORTS       NAMES
33ecc4cc2bfb   yolo-colo/scan    "/docker-entrypoint.…"   21 hours ago    Up 21 hours    25/tcp      scan_1
8b86f724e188   yolo-colo/rdns    "/bin/sh -c /entrypo…"   21 hours ago    Up 21 hours    53/udp      rdns_1
15d771655355   yolo-colo/rdns    "/bin/sh -c /entrypo…"   21 hours ago    Up 21 hours    53/udp      adns_1
```

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

supported by:
RIPE NCC Community Projects Fund
LWLCOM connecting you.
OPENFACTORY
Virtuacloud
DE-CIX
WIIT THE PREMIUM CLOUD

Running Reliable Network Measurements

```
root@msrmnt.example.com:/opt/yolo-colo# docker compose up -d
root@msrmnt.example.com:/opt/yolo-colo# docker ps
CONTAINER ID   IMAGE            COMMAND                  CREATED         STATUS         PORTS       NAMES
33ecc4cc2bfb   yolo-colo/scan   "/docker-entrypoint.…"   21 hours ago    Up 21 hours    25/tcp      scan_1
8b86f724e188   yolo-colo/rdns   "/bin/sh -c /entrypo…"   21 hours ago    Up 21 hours    53/udp      rdns_1
15d771655355   yolo-colo/rdns   "/bin/sh -c /entrypo…"   21 hours ago    Up 21 hours    53/udp      adns_1
```

- No TCP for DNS
- Maybe rdns_1 stops resolving
- Maybe DNS is not delegated to adns_1
- Maybe rdns_1 just stubs vs. q1/q8/q9 and we do not know if queries we see at adns_1 just come from… us…

© Constanze Dietrich

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Security Measurements

- You want to do some SSH related scanning on the Internet
- Your colleague forces you to figure out why a small fraction of hosts in his AS you use for testing are not found

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

supported by:
RIPE NCC Community Projects Fund
LWLCOM connecting you.
OPENFACTORY
Virtuacloud
DE-CIX
WIIT THE PREMIUM CLOUD

Security Measurements

- You want to do some SSH related scanning on the Internet
- Your colleague forces you to figure out why a small fraction of hosts in his AS you use for testing are not found
- All packets go out, replies go out at the targets, but never arrive

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Security Measurements

- You want to do some SSH related scanning on the Internet
- Your colleague forces you to figure out why a small fraction of hosts in his AS you use for testing are not found
- All packets go out, replies go out at the targets, but never arrive

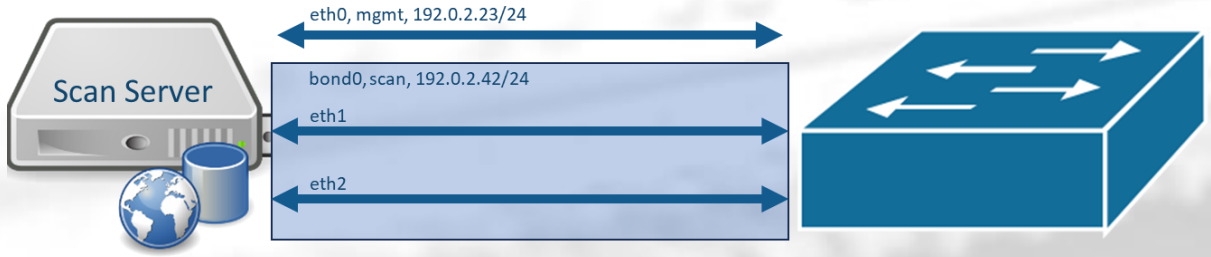

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Security Measurements

- You want to do some SSH related scanning on the Internet
- Your colleague forces you to figure out why a small fraction of hosts in his AS you use for testing are not found
- All packets go out, replies go out at the targets, but never arrive

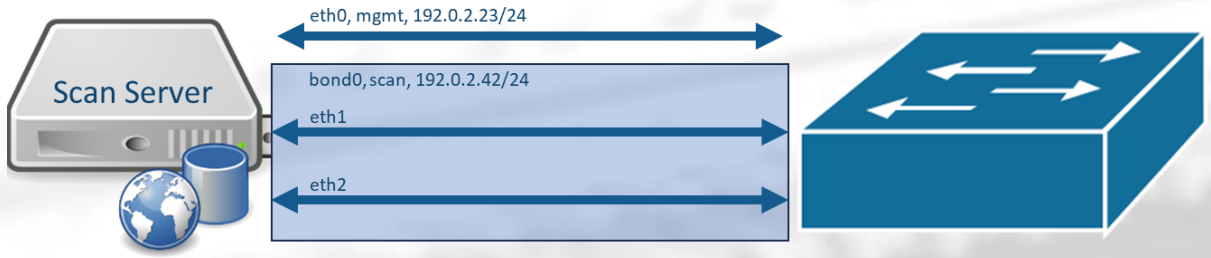

Rules:
from 192.0.2.23/24 lookup table 101
from 192.0.2.42/24 lookup table 102

Tables:
101: 0.0.0.0/0 via 192.0.2.1 dev eth0
102: 0.0.0.0/0 via 192.0.2.1 dev bond0

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Making people *really* unhappy

- ```
  sshd[8423]: Failed password for user root from 192.0.2.42 port 44216 ssh2
  sshd[8424]: Failed password for user root from 192.0.2.42 port 19561 ssh2
  sshd[8425]: Failed password for user root from 192.0.2.42 port 39148 ssh2
  …
  ```
- We are not *actually* attempting authentication, though…
- Still a great opportunity to make many friends via the (partially anonymous) chat service under abuse@
- Even greater way to get the org's whole prefix blocked
- Amazing way to find hidden middleboxes in the org's network

Tobias Fiebig – tfiebig@mpi-inf.mpg.de
  ```

The Internet is Complex & Lore Driven

- Understanding how the Internet works is difficult
  - "So, with whom do you peer, and with whom are you *actually* peering?
- "All my friends are construction companies"
  - And I upstream at least:
    - 2x ISP
    - 1x Insurance Company / Finance Company / Bank
    - 1x Media/Publishing
- 'Peering agreements are generally established between the legal departments of two corporations and then communicated to IT for implementation.'

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Network Protocols are MORE Complex

- "Whoever put the 'S' for simple into SMTP & SNMP, also put the 'L' for lightweight into LDAP."
- Asking a "widely used DNS server's" dev (10y+), whether it would be feasible for a person to write DNS software after ~1y of working on/with DNS that can be safely run against the Internet *without* breaking anything:

"Well, I wouldn't even trust myself to write some new DNS software that does not break anything when thrown against the whole Internet…"

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Writing Reliable Measurement Software

- The Internet is full of corner cases: Account for all!
- Be aware of all (unwritten) rules of your protocol of choice.
- Implement a *reliable* measurement tool, ideally reusing as much existing (tested) software as possible
- Be, in general, a good and experienced programmer able to write software able to interact with all systems on the Internet (not breaking them even if you just do standard-compliant things when interacting with them)
- Version Control, tests… proper development!

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Running Reliable Network Measurements

- Be an experienced SysOp
- Know about all the things involved (and available tools!)
- Monitor your stack
  - Historic for bottlenecks (you may just measure your IOPS)
  - Real-time for reliability
- Have an end-to-end understanding
- Make sure the setup is self-contained

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

supported by:
RIPE NCC Community Projects Fund
LWLCOM connecting you.
OPENFACTORY
Virtuacloud
DE-CIX
WIIT THE PREMIUM CLOUD

Running Ethical Measurements

- Consider all possible unintended harms
  - "Yeah, we know, the Internet is made from duct tape and bubble gum, and this would be an issue; So we just don't talk about it!"
- Get ethics approval
- Do probe attribution
  - rDNS, RIR Whois, running webserver etcetcetc.
- Handle 24/7 opt-out and abuse
- Have a maintained block-list

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

The PhD we Need

- Thoroughly understand the protocol stack they are measuring, including operational lore and lived experience since the inception of these protocols
- Be versed in the domain of available implementations to identify components they can use to construct their measurement setup
- Be experienced programmers and versed in software development in general to follow development best practices and produce tested and reliable code
- Be experienced system administrators–or have such institutional support–to setup the measurement system, including all basic services the system depends on, including historic and real-time monitoring of all components

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

The PhD we Need

- Thoroughly understand the protocol stack they are measuring, including operational lore and lived experience since the inception of these protocols
- Be versed in the domain of available implementations to identify components they can use to compose their measurement setup
- Be experienced programmers and versed in software development in general to follow development best practices and produce tested and reliable code
- Be experienced system administrators–or have such institutional support–to setup the measurement system, including all basic services the system depends on, including historic and real-time monitoring of all components

This is not a PhD student;

This is a whole IT department.

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

supported by: RIPE NCC Community Projects Fund · LWLCOM connecting you. · OPENFACTORY · Virtuacloud · DE-CIX · WIIT THE PREMIUM CLOUD

The Reality of a PhD

- 4-8 Years
- ~4 'Top-Tier' papers
- New research advancing the field
- Embedded in related work
  (Meaning: You have to read it!)
- Joining after a bachelor's degree (US)
  or master's degree (most-other-ish)

- **First paper should be under submission after ~1 year!**
- **People under pressure will do people things**

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

LPU: Least Publishable Unit

- It is easier to publish what is ‘novel’, ‘flashy’, and ‘has impact’
- The LPU of “we spent six months on ‘engineering’, which was really challenging, but did not change our results” is *zero*
- Academia is *very* competitive:
  - You are measured on #papers, #supervised students, #acquired grants, #committee tasks, #courses, #…
  - There is always someone better than you in everything at once
  - You need more papers, and must diversify your research interests
- Some things you just learn by doing, especially running infrastructure and things on the Internet

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Changing the Game: Building measurement.network

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

supported by:
RIPE NCC Community Projects Fund
LWLCOM connecting you.
OPENFACTORY
Virtuacloud
DE-CIX
WIIT THE PREMIUM CLOUD

Core Idea Behind measurement.network

- Have available infrastructure to run measurements from
  - Not 'tied' to any organization with publishing 'incentive'
  - Well-known and blockable
  - Taking ops basics off the plate of researchers (monitoring, base infrastructure, (r)DNS, ensuring unfiltered PPS)
  - Make more™ things accessible (LIR services, resources etc.)
  - Support other science processes (open data, reviewing etc.)
- Loop in people who do things in practice to review and guide measurements / research (before unhappy XYZ-NOG ML threads have to be started)

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

The Progress

- ✓Get AS: AS211286 (main) & AS215250 (V4LESS-AS)
- ✓Get IP: 141.39.220.0/22 & 2a0d:8d04::/32
- ✓Get upstream: AS50629, AS58299, AS59645, Community-IX
- ~Get routers: Juniper Routers for DUS and BER (deploy TBD)
- ✓Get servers: Co-use of AS59645 in DUS/BER, cluster in SBR
- ~Setup things
- ~Run services / measurements
- ❑Do reviews

Tobias Fiebig – tfiebig@mpi-inf.mpg.de

Key Take-Aways

- Academics try their best to do good work
- Realities of academia can stand in the way
- There is (still) far too little interaction between academia and people running systems

- `measurement.network` is a fun project aimed at making research more reliable, accessible, and tied in practice
  - Consider to review/mentor & contribute:
    Write to contact@measurement.network

Tobias Fiebig – tfiebig@mpi-inf.mpg.de