



measurement.network: Organizing Academic Network Measurements

Tobias Fiebig
Max-Planck Institut für Informatik

supported by:



RIPE NCC
Community Projects Fund



LWL.COM
connecting you.



OPENFACTORY

Virtuacloud

What is wrong with academics?



Spoiler alert:
There will be a talk at RIPE88 on this



Network Measurement



- 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

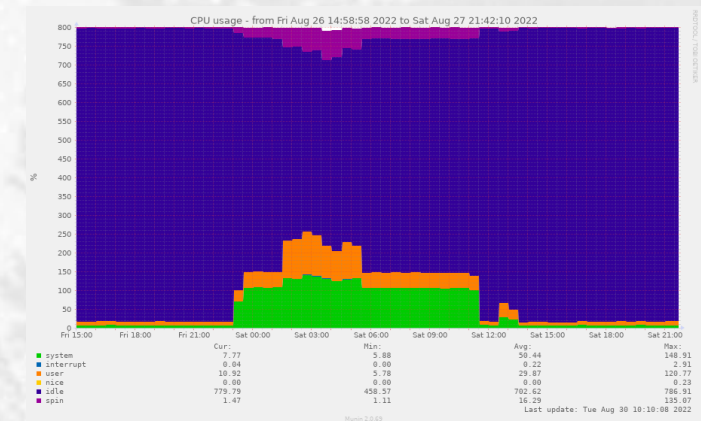

Writing Reliable Measurement Software



Writing Reliable Measurement Software



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



Writing Reliable Measurement Software



No SMTP



Current Status:

Status Information:

CRITICAL (for 0d 0h 0m 7s)

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

Writing Reliable Measurement Software



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



```

88 54126 table-mysql GIO   fd 3 read 92 bytes
89      "X\0\0\0
90      5.5.5-10.6.7-MariaDB\0%\M- [...] \
      mysql_native_password\0"
91 54126 table-mysql RET   rcvfrom 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\0\0opensmtpd\0 [...] \
      smtpd\0mysql_native_password\01\^C_os\ \
      OpenBSD\ f client_name
95      libmariadb\^D_pid\^E54126\^0 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



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.’



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...”



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



The PhD we Need



- Thoroughly understand the protocol stack they are measuring, including operational lore and lived experience since the inception of these protocols

**This is not a PhD student;
This is a whole IT department.**

- 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 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





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!**

MARRIAGE vs. The Ph.D.

	 Marriage	 Ph.D.
Typical Length:	7.5 years	7 years
Begins with:	A proposal	A thesis proposal
Culminates in a ceremony where you walk down an aisle dressed in a gown:	✓	✓
Usually entered into by:	Foolish young people in love	Foolish young people without a job
50% end in:	Bitter divorce	Bitter remorse
Involves exchange of:	Vows	Know-how
Until death do you part?	If you're lucky	If you're lazy

JORGE CHAM © 2010
WWW.PHDCOMICS.COM

Source: "Piled Higher and Deeper" by Jorge Cham
<https://phdcomics.com/comics/archive.php?comid=1296>

Operational Support



- IT departments dislike measurement setups in their network
- Some researchers scan from public clouds (hard to block)
- Ethics committees / IRBs do not understand lore & technology
- Middleboxes everywhere
- Infrastructure often does not survive PhD students' departure





Changing the Game: Building measurement.network

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



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 (before unhappy XYZ-NOG ML threads have to be started)



The Plan



- Get AS
- Get IP
- Get upstream
- Get routers
- Get servers
- Setup things
- Run services / measurements
- Do reviews



The Plan



- ✓ 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



Setup Things



- Core Infrastructure is running (network, DNS, web, mail, ...)
- Router in DUS deployed, BER pending (currently SW VyOS)
- Rack in SBR with 1TB mem Virt. Cluster
- Infrastructure controlled with Puppet + Self-written config generator for VyOS/JunOS
- Setup auth. system (Keycloak) for reviewers & researchers
- Patched submission system (HotCRP) to work with auth. system





- Running public services:
 - <https://bttf-whois.measurement.network/>
 - <https://www.email-security-scans.org/>
 - <https://v6only-resolver.measurement.network/>
 - V4LESS-AS (more later)
 - Artifact Evaluation (more later)
- Hosted measurements so far:
 - Recursive DNS measurements for the MSc thesis of a student from JP
 - Active DNS measurements for a project of one of my students (eat your own dogfood...)



Run services: V4LESS-AS



- <https://measurement.network/services/v4less-as/>
 - Test eBGP for IPv4 routes without v4 addresses on IFs (RFC8950)
 - Less for academics, more for operators
- Present at IXPs: BCIX, FNC-IX, FogIXP, France-IX (Lille&Paris)
- Has passive sessions for all IX members pre-configured
- Hosts test-targets for various cases (MTU 1400/1500, with/without GUA v4 ICMP source for on-path routers)
 - 4x RIPE Atlas Probe
 - 4x NLNOG Ring Node



Run services: Artifact Evaluation



- AE: Check code/data of (to be) published papers
- Provide authors & reviewers with clean VMs to run code on and measurements from
 - <https://measurement.network/review-support/artifact-evaluation/>
- Wrote OSS integration to VM backend (Proxmox) for HotCRP
- Two conferences so far:
 - PETS (Privacy Enhancing Technologies Symposium) 2024.2 & 2024.3
 - ACSAC (Annual Computer Security Applications Conference) '24





Next Steps & Todo



Next Steps: Todo



Operational

- Get more routers deployed
- Maybe get a rack in FMT
- Ticket system for abuse@
- Add storage to host collected measurement data
- Find people willing to join the operations part

Organizational

- Write policy documents
 - Privacy policy
 - Statutes
 - Review
- Find reviewers/mentors
- Start reviewing things and run measurements



Next Steps: Tolong-termndo



Operational

- Add anycast setup
- Find (a bit) more v4 ;-)
- Add new virt. systems
- Add GPUs and ARM to artifact evaluation system

Organizational

- Find a body who could pick up the governance of this
- Find good way to create a legal entity around it (if not provided by governance body) and LIR-out



Key Take-Aways



- Things™ are happening; The project is growing and starts being actively useful for the community
 - Public Services & Artifact Evaluation
 - First measurements
- Governance is the difficult part™
- Generally on-track
- If you are at an IXP with V4LESS-AS, setup an RFC8950 session
- If you are an IXP, think about sponsoring a V4LESS-AS node
- Consider to review/mentor & contribute:
Write to contact@measurement.network

- ✓ 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: There is always more to do
- ~ Run services / measurements: Some things are already there
- Do reviews: This is the next big step!

