enum.at
net.communications

crawler.enum.at

RIPE56, May 2008, Berlin, Germany

Alexander Mayrhofer, enum.at GmbH

What‘s the ENUM crawler?

enum.at
net.communications

- Explore the e164.arpa namespace, and find all numbers with NAPTRS
- Why?
  - FAQ #1: „How many people can i reach using this ENUM thingy?“
  - „Good question. Let‘s find out.“
  - Plus, it‘s always fun to play with PostgreSQL

enum.at ENUM crawler

searching e164.arpa for NAPTR sets

This is a DNS-based crawler which crawls trough e164.arpa, the top level domain designated for ENUM. Specifically, it looks for NAPTR resource records sets, and tries to discover the whole "golden" ENUM tree

Status: **585876** ENUM numbers (containing **674372** NAPTR records). Crawling speed: **272** numbers in the last minute.

What is the ENUM crawler?

The ENUM crawler uses certain properties of the e164.arpa namespace to discover all numbers for which ENUM entries (NAPTRs) exist. A full crawling round takes a little longer than a week, currently, so that new numbers should appear after approximately that delay. However, certain nameserver implementations might prevent the crawler to discover all enum-enabled numbers.

top country codes

#	country name	E.164	ENUMs
1.)	Romania	+40	410674
2.)	United Kingdom	+44	51001
3.)	Austria	+43	32156
4.)	Poland	+48	25032
5.)	Germany	+49	22434
6.)	Sweden	+46	14234
7.)	Korea, Republic of	+82	13114
8.)	Japan	+81	10030
9.)	Czech Republic	+420	3168
10.)	China	+86	1657
11.)	Norway	+47	1157
12.)	Slovakia	+421	723
13.)	United Arab Emirates	+971	273
14.)	Australia	+61	121
15.)	Ireland	+353	50
16.)	Indonesia	+62	33
17.)	Singapore	+65	8
18.)	Finland	+358	3
19.)	Greece	+30	2
20.)	Netherlands	+31	2
21.)	Liechtenstein	+423	2
22.)	Brazil	+55	2

recently discovered

+971 504581166 (9 hours ago, 1 NAPTR)
+43 1599664372039 (9 hours ago, 1 NAPTR)
+43 1599664374299 (9 hours ago, 1 NAPTR)
+43 159966437203 (9 hours ago, 2 NAPTRs)
+43 1599664379279 (9 hours ago, 1 NAPTR)
+43 159966437927 (9 hours ago, 3 NAPTRs)
+43 159966437429 (9 hours ago, 2 NAPTRs)
+86 13588886846 (9 hours ago, 1 NAPTR)
+971 504463334 (12 hours ago, 4 NAPTRs)
+971 504584554 (12 hours ago, 1 NAPTR)

How does it work?

enum.at
net.communications

- Namespace easy to crawl
  - 0..1..2..3..4..5..6..7..8..9
- Still, a huge namespace (10^15)
  - 232830 times larger than the IPv4 space
  - Lucky enough, 340282366920938463463374 times smaller than IPv6 ☺
- To the rescue: DNS empty non-terminals
  - Original idea stolen from Peter Koch
  - RFC 4592 mandates servers to return „NOERROR“ if there‘s something „below“ an empty label

Empty Non-Terminals

enum.at
net.communications

Zone example:

```
1.2.3.4   IN   NAPTR   (recordsX)
  3.3.3   IN   NAPTR   (recordsY)
5.5.2.5   IN   NS      some.where.else.
```

Query examples:

```
  1.2.3.4    NAPTR? -> NOERROR  (recordsX)
    3.3.3    NAPTR? -> NOERROR  (recordsY)
    7.7.7    NAPTR? -> NXDOMAIN
    2.3.4    NAPTR? -> NOERROR  (empty)
    5.2.5    NAPTR? -> NOERROR  (empty)
  7.5.2.5    NAPTR? -> NXDOMAIN
```

Use in the ENUM crawler

enum.at
net.communications

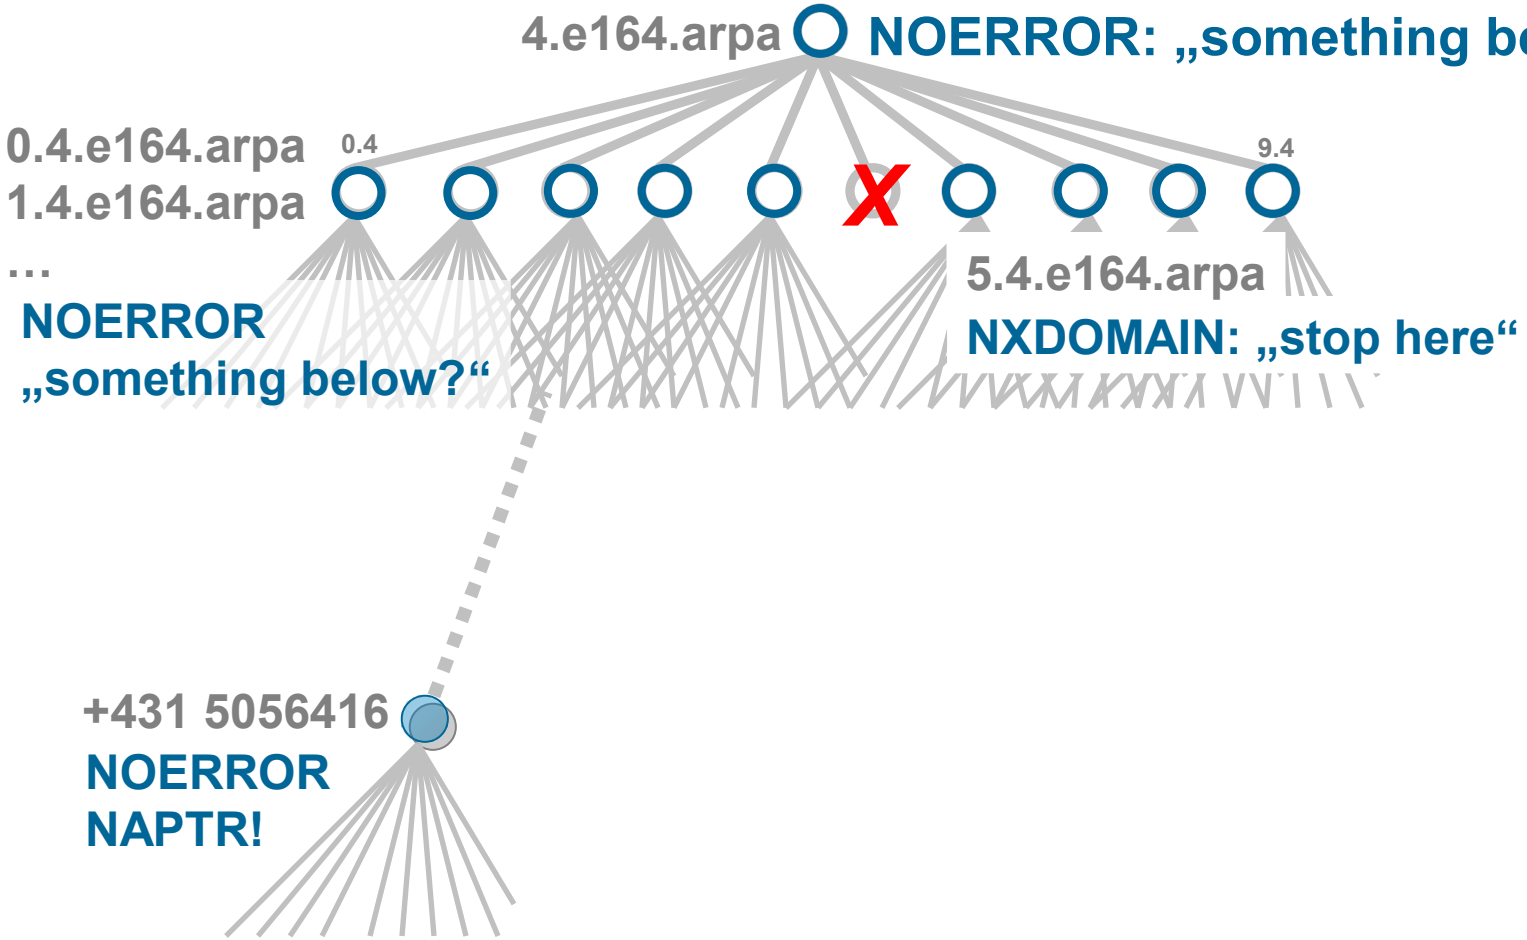

Issues with the approach

enum.at
net.communications

- Some nameserver still return(ed) NXDOMAIN for empty non-terminals
  - Clarified in RFC 4592
  - ENUM space below not discovered
  - Crawler deletes existing information „below“
  - Older BIND, some PowerDNS versions
- One such case was one of the e164.arpa secondaries :-/
  - (temporary) loss of 4.e164.arpa in the crawler
  - Safeguard now: „never delete more than 1000 nodes“

What the crawler stores

enum.at
net.communications

- Number status
  - „X“: ENUM domain does not exist („stop“)
  - „E“: Empty („something below“)
  - „N“: NAPTRs found
  - „W“: Wildcards found
    - Trick: re-query with _blafusel.3.4.e164.arpa
  - „P“: Problem (Lame delegation, timeout)
  - „L“: Locked
    - „Creative“ numberspaces, don‘t query
- Full NAPTRs (if any)
- DNS Query time
- Discovered, last queried timestamps

Some fun results

enum.at
net.communications

- http://crawler.enum.at/
- 585876 Numbers with 674372 NAPTRs
  - 568215 „normal“, 17658 „wildcarded“
- 22 countries
  - Top: .ro: 410671, .uk: 51001, .at: 32156
- Number length:
  - Shortest with NAPTR: +40 778
  - Lots of numbers with 15 digits (crawler stops)
  - Average: „normal“ 11.25, „wildcard“ 11.5 (??)

Popular Enumservices

enum.at
net.communications

count	lowercased
414296	**sip+e2u**
58725	e2u+sip
34010	e2u+tel
6017	**e2u+ep**
6015	**e2u+fixed**
2858	e2u+ifax:mailto
2746	e2u+sms:mailto
2720	e2u+http
824	e2u+email:mailto
821	e2u+mailto
741	tel+e2u
480	e2u+h323
436	http+e2u
408	e2u+fax
348	web:http+e2u
317	mailto+e2u
302	e2u+iax2
300	e2u+voice:sip
223	e2u+x-skype:callto
200	fax+e2u
109	e2u+web:http
105	e2u+email
104	e2u+msg
79	e2u+voice:tel
66	voice:sip+e2u
66	e2u+h323:voice
61	e2u+iax
56	e2u+msg:mailto
43	e2u+mobile
33	iax2+e2u
31	e2u+voice
22	e2u+web
20	e2u+mail:mailto
19	e2u+icq
17	e2u+fax:tel
11	e2u+pstn:tel
11	e2u+sms:tel
11	e2u+web:https
10	e2u+iax:iax2
10	e2u+vcard
7	e2u+vcard:plain
7	e2u+sms
7	e2u+mms:tel
4	e2u+ftp

Rare Enumservices ☺

enum.at
net.communications

count	lower
1	e2u+address
1	smtp+e2u
1	e2u+smp
1	e2u+
1	address+e2u
1	e2u+ems:tel
1	e2u+ft:ftp
1	e2u+voice:h323
1	email+e2u
1	e2u+im
1	e2u+mailto:msg
1	sms:tel+e2u
1	e2u
1	e2u+smtp
1	message:http+e2u
1	e2u+info
2	e2u+voip
2	e2u+vpim:mailto
2	e2u+vpim:ldap
2	e2u+mail
2	e2u+domainkey
2	e2u+www
2	e2u+email:email
2	e2u+void:mailto
2	sip+d2u
2	e2u+ifax
2	e2u+paging
2	sms+e2u
3	e2u+relay
3	e2u+u
3	e2u+xmpp
3	e2u+pres
3	e2u+ldap
4	e2u+ftp
4	e2u+mms:mailto
7	e2u+vcard:plain
7	e2u+sms
7	e2u+mms:tel
10	e2u+iax:iax2
10	e2u+vcard
11	e2u+pstn:tel
11	e2u+sms:tel
11	e2u+web:https
17	e2u+fax:tel

Other observations

enum.at
net.communications

- Average query time: 86ms overall, 65ms for numbers with NAPTRs.
  - From Server in Austria
- Just _one_ (1) „sips:“ ENUM entry
  - sips:lendl@nic.at43.at …
- 99.99494% of NAPTRs use „!“ as regexp delimiter (some use „space“ ??)
- 16 of the top 20 „ENUM-countries“ have the color red in their flag.

Software / Hardware

enum.at
net.communications

- Stone-age dual CPU server
- Debian Linux
- PostgreSQL 8.1 database
  - ~8 million rows
  - Some queries are … challenging.
- Crawler + Web Frontend in PHP
  - Net_DNS module
  - Blueprint CSS framework
  - Cute flag icons from famfamfam.com
  - Smarty template engine

Questions?

enum.at
net.communications

Thank you!

Alexander Mayrhofer
Alexander.mayrhofer@enum.at