

The DNS Scare of 2008

What was it all about and how did
Austria react?

Otmar Lendl <lendl@cert.at>


Otmar Lendl
18. September 2008


- The bug
 - Background on DNS forgeries
 - What was new
- Austria reacts to VU#800113
 - Patch statistics
- What are the long-term solutions?
 - Various ideas

- Global, distributed Database
- Input: Domain name
- Output: Resource Records
 - A, AAAA IP addresses
 - MX Mail Exchanges
 - CNAME Canonical Name
 - NS Name Servers
 - PTR Reverse Lookup
 - SRV, TXT, DS, NSEC, NSEC3
- Transport: mainly UDP
- Lot's of caching

You all know that already ...

- Authoritative Nameservers
- Recursive Nameservers
- Stub Resolvers





```
$ dig cert.at MX
[...]
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 8861
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 1

;; QUESTION SECTION:
cert.at.                IN      MX

;; ANSWER SECTION:
cert.at.                7200   IN      MX      100 nuwen.cert.at.

;; AUTHORITY SECTION:
cert.at.                6450   IN      NS      ns1.cert.at.
cert.at.                6450   IN      NS      ns5.univie.ac.at.

;; ADDITIONAL SECTION:
nuwen.cert.at.         7200   IN      A       83.136.33.135
```

This is the safety net

```
$ dig cert.at MX
[...]
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 8861
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 1

;; QUESTION SECTION:
;cert.at.                IN      MX
```

```
;; ANSWER SECTION:
cert.at. 600 IN      MX      10.0.0.1
;; AUTHORITY SECTION:
cert.at. 600 IN      NS      ns1.univie.ac.at.
cert.at. 600 IN      NS      ns2.univie.ac.at.
;; ADDITIONAL SECTION:
nwwen.cert.at. 7200 IN      A       83.136.33.135
```

16 bits of entropy might have been enough in 1987, but not in 2008.

- Legitimate Servers can play games, too.

```
;; ADDITIONAL SECTION:
nuwen.cert.at.          7200    IN      A       83.136.33.135
```

- What stops a server from adding?

```
;; ADDITIONAL SECTION:
nuwen.cert.at.          7200    IN      A       83.136.33.135
xxx.                    50000   IN      NS      ns1.cert.at.
xxx.                    50000   IN      NS      ns2.cert.at.
```

- “in-bailiwick checks”: only accept what you were asking for. See RFC 2181 for details.

- State of the game in 2008
- An attack needs to match
 1. The question section of the reply packet is equivalent to that of a question packet currently waiting for a response
 2. The ID field of the reply packet matches that of the question packet
 3. The response comes from the same network address the question was sent to
 4. The response comes in on the same network address, including port number, as the question was sent from
- Calculations on forgery success probabilities

- A good PRNG is needed
 - Older Bind versions simply incremented
 - Good entropy isn't that easy
- Avoid multiple outstanding requests for the same question
 - n queries * m forgeries give good odds for the attacker.

- How often can an attack take place?
 - Forcing a query helps the attacker.
 - One attack per TTL

- “The calculations above indicate the relative ease with which DNS data can be spoofed. For example, using the formula derived earlier on an RRSet with a 3600 second TTL, an attacker sending 7000 fake response packets/s (a rate of 4.5Mb/s), stands a 10% chance of spoofing a record in the first 24 hours, which rises to 50% after a week.”

- Entropy from ID (16) and server (2) is barely enough.

- Recommendations:
 - Use Source Port Randomization to get another 15 bits.
 - Make sure you don't mess up the entropy

- Dire Warning
 - Insufficient entropy in ID
 - Multiple outstanding requests
 - Fixed source port
- Recommendation
 - Update
 - Implement Source Port Randomization
 - Restrict Recursion
 - Filter spoofed IP traffic

- RFC 3833, 2.3 proved to be right
 - “Perhaps the most interesting class of DNS-specific threats are the name chaining attacks. These are a subset of a larger class of name-based attacks, sometimes called "cache poisoning" attacks.”
- The TTL wall had been broken.
 - Idea: don't attack the Domain itself, go for random hostnames within the target domain.
 - A way to initiate recursion is needed.

;; QUESTION SECTION:

;345678.example.org. IN A

;; ANSWER SECTION:

345678.example.org. 3600 IN A 192.0.2.1

;; AUTHORITY SECTION:

example.org. 3600 IN NS evil1.example.net.

example.org. 3600 IN NS evil2.example.net.

Source: IETF namedroppers list. (P. Koch, T. Finch)

;; QUESTION SECTION:

345678.example.org. IN A

;; ANSWER SECTION:

345678.example.org. A 192.0.2.1

;; AUTHORITY SECTION:

example.org. NS 345678.example.org.

;; QUESTION SECTION:

;345678.www.example.org. A

;; AUTHORITY SECTION:

www.example.org. NS evil1.example.net.

www.example.org. NS evil2.example.net.

```
;; QUESTION SECTION:
```

```
;345678.www.example.org.          A
```

```
;; AUTHORITY SECTION:
```

```
www.example.org.          NS      evil.example.net.
```

```
;; ADDITIONAL SECTION:
```

```
evil.example.net.        A      192.0.2.53
```

```
;; QUESTION SECTION:
```

```
;345678.example.org.          IN      A
```

```
;; ANSWER SECTION:
```


```
345678.example.org.         CNAME   www.example.org.  
www.example.org.           A       192.0.2.80   ; evil
```

- Man-in-the-middle *everything*
 - Phishing
 - Email hijacking
- DoS
- Password reset emails
- What about SSL?
 - CA email-loop
 - CA whois lookup

- Warnings on all channels
- Wait a second ...
 - the patch changes the query pattern
 - resolver in Austria query for domains under .at
 - we have query-logs of the .at nameservers
 - important resolvers ask for a lot of domains

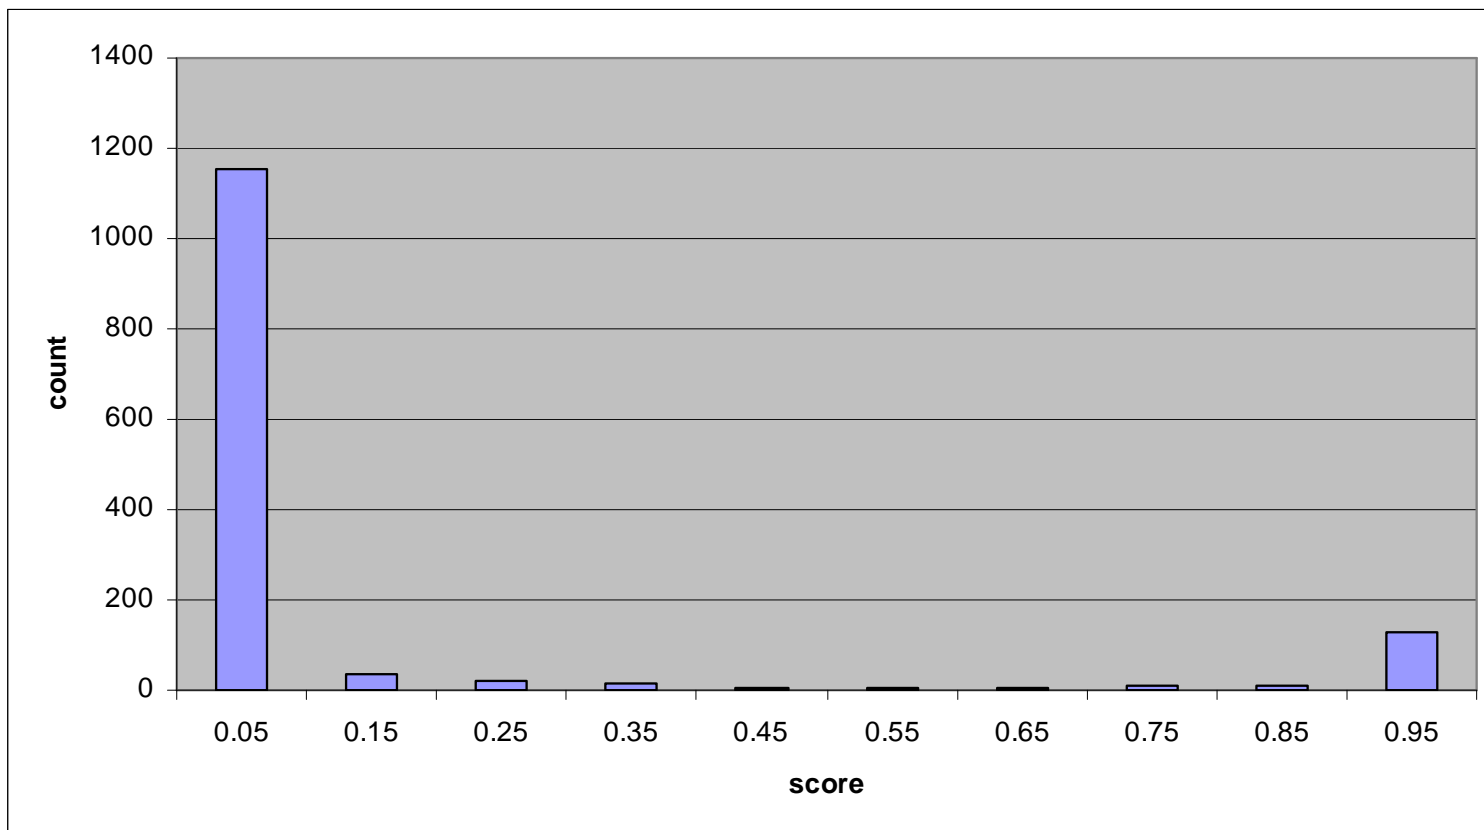
We can monitor the patching progress

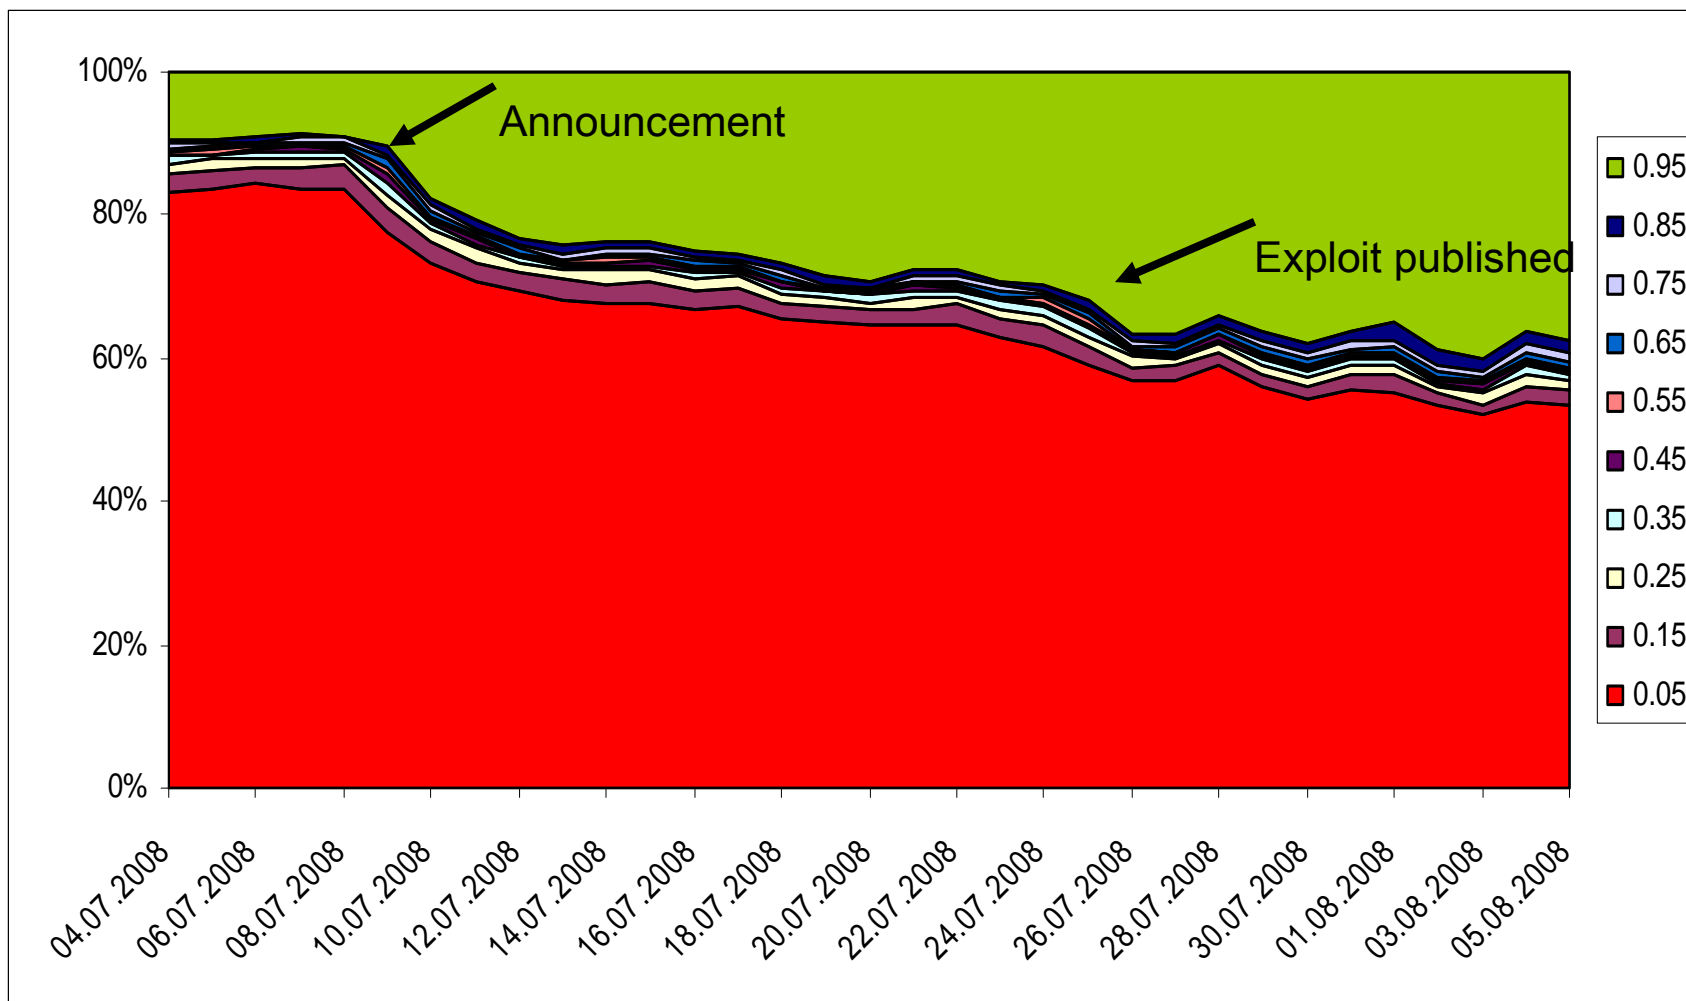
$$score = \frac{portchanges}{queries} * \frac{ports}{min(queries, 65536)}$$

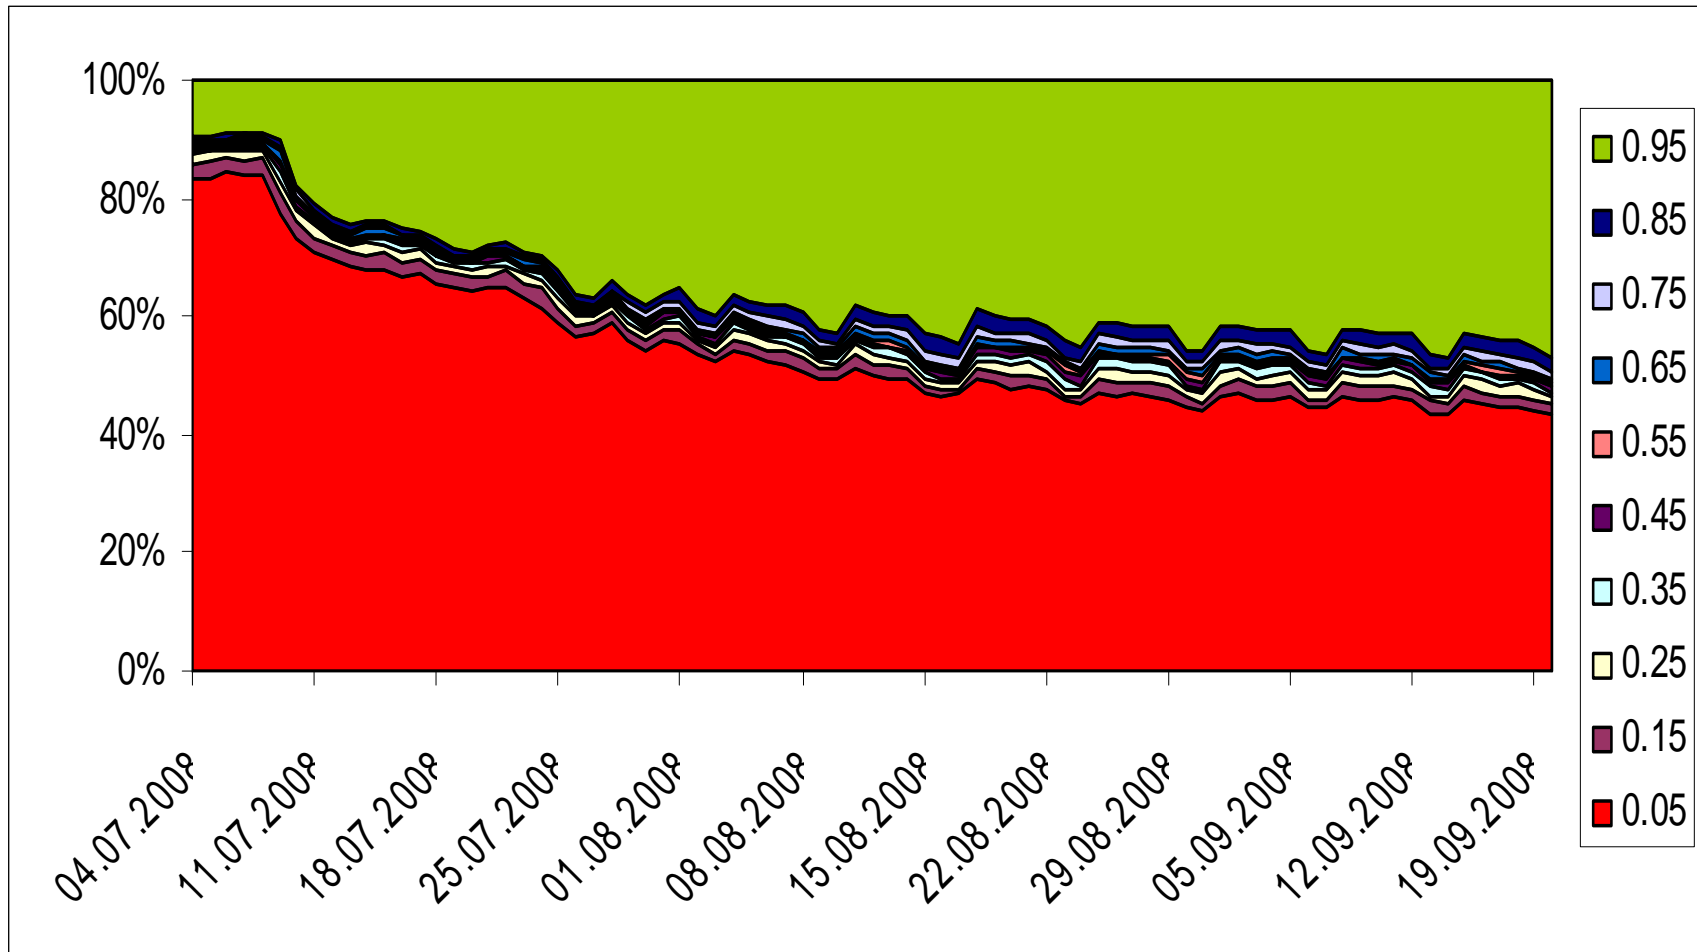



- Query log of one server in Austria
- ~ 10M queries / day (from Austrian IP space)
- from ~ 13.000 IP addresses
- ~2000 of which sent more than 100 queries
- Known domain catchers / registrars excluded
- Data starts with July 4th, i.e. before the announcements


Thanks to Markus Heimhilcher from the .at nameserver management team!











Looks like a number of ISPs reworked their DNS setup over the summer.

- Source Port Randomization is a stopgap.
- Brings the time to exploit down from “oh sh*t, we’re f***ed” to “not really secure”
- So: what is the outlook? What can be done?

1. Get more entropy
 2. Changing resolver behavior
 3. Cooperation from both sides
 4. Changing the protocol
 5. Agile defenses
-
- See namedroppers list archive
 - [draft-wijngaards-dnsexst-resolver-side-mitigation-00](#)

- QID
- Source port
- Source Address
 - IPv6!
- Destination Address
- 0x20
 - Case is preserved, but not relevant

- Repeated Queries
 - Same server
 - Different servers
- Use TCP
- Time spaced
- New rules on how to trust information
 - Never cache glue
 - Be careful when overwriting the cache
 - i.e. update the RFC 2181 rules

- TSIG
- SIG(0)
- IPsec

- DNSSEC

- More Entropy
 - EDNS ECHO
 - draft-hubert-ulevitch-edns-ping-00
 - Query name hacks
 - Prefixes / Postfixes
 - <http://www.jhsoft.com/dns-xqid.htm>
 - <http://www.ops.ietf.org/lists/namedroppers/namedroppers.2008/msg01555.html>
 - QCount > 1
 - DNS Cookies
 - draft-eastlake-dnsexp-cookies
- More Crypto
 - <http://dnscurve.org/>

All need protection against down-grade attacks

- Some of the mitigation strategies carry a cost, thus:
 1. Try to notice an attack
 - Many mismatches
 - Timing
 2. React
 - Refuse to cache
 - Fallback to TCP
 - Multiple queries

- The DNS Scare was justified.
- Patching is progressing slowly.
- We're far from a long-term solution
 - Low-impact patches?
 - DNSSEC?

Questions ?