TECHNISCHE UNIVERSITÄT DRESDEN

DRESDEN concept

Thorsten Strufe

Resilient Networking

Module 5: Name Resolution / DNS

Disclaimer: This module prepared in cooperation with Mathias Fischer, Michael Roßberg, and Günter Schäfer

Dresden, SS 19

Module Outline

Overview of DNS

Known attacks on DNS
- Denial-of-Service
- Cache Poisoning

Securing DNS
- Split-horizon DNS
- DNS Cookies (RFC 7873)
- DNSSEC
- DNSCurve
- PNRP
- GNS

DNS – The Domain Name System

PRIVACY AND SECURITY

What is DNS?

Naming Service for (almost all) Internet traffic

Lookup of (resolve)

- Host-Addresses
- Mail-Servers
- Alias Names
- Alternative Name Servers
- ...

Distributed Database consisting

of multitude of servers

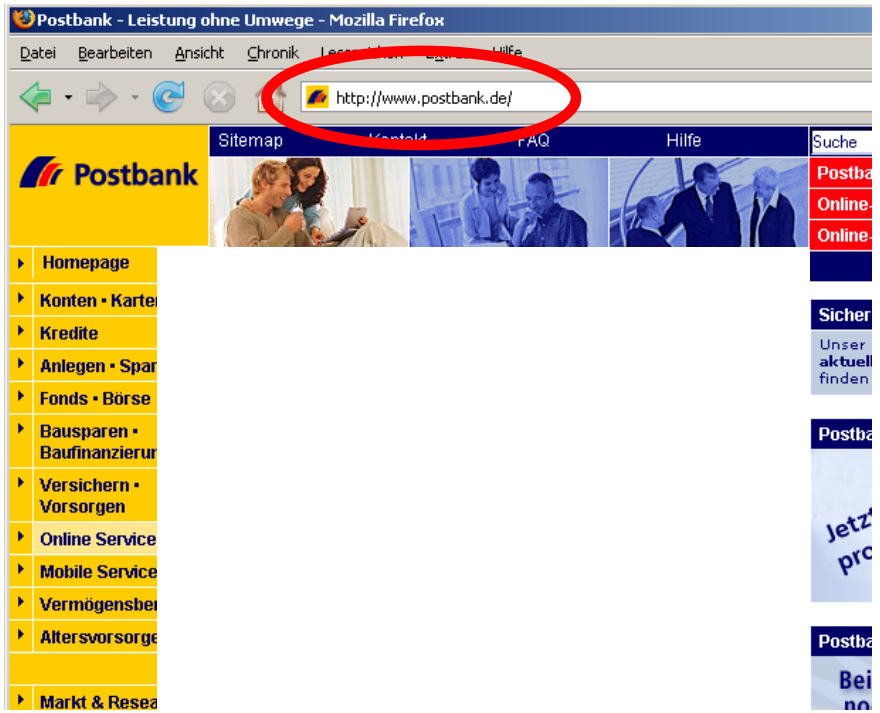

DNS – Names

People: many identifiers:

- SSN, name, passport #

Internet hosts, routers:

- IP address (32 bit) - used for addressing datagrams
- "Name", e.g., www.yahoo.com - used by humans

Q: Map between IP addresses and name ?

Domain Name System:

Distributed database implemented in hierarchy of many *name servers*

Application-layer protocol: hosts, routers, name servers communicate to *resolve* names (address/name translation)

- Note: core Internet function, implemented as application-layer protocol
- Complexity at network's "edge"

DNS – what does it do?

PRIVACY AND SECURITY

DNS services

- Hostname to IP address translation
- Host aliasing
  - Canonical and alias names
- Mail server aliasing
- Load distribution
  - Replicated Web servers: set of IP addresses for one canonical name

Why not centralize DNS?

- Single point of failure
- Traffic volume
- Distant centralized database
- Maintenance
- *does not scale!*

What does this „it scales“ mean anyways!?

DNS – Data Organization: Domains / Zones

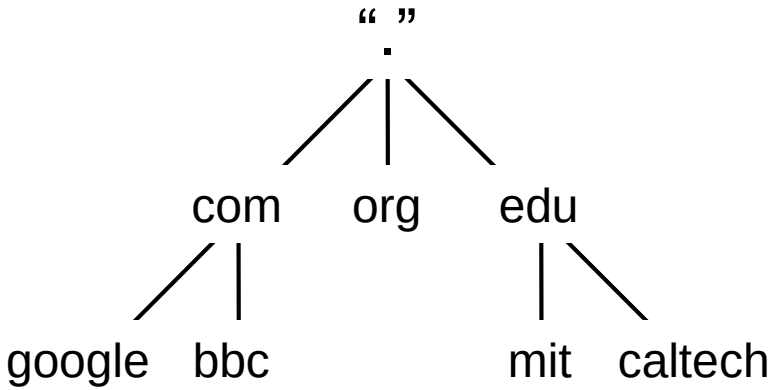

- Structured Namespace
- Hierarchical organization in sub domains/zones
- Sourced at “root zone” (“.”)
- Parent zones maintain pointers to child zones (*“zone cuts”*)
- Zone data is stored as “Resource Records” (RR)

Distributed, Hierarchical Database

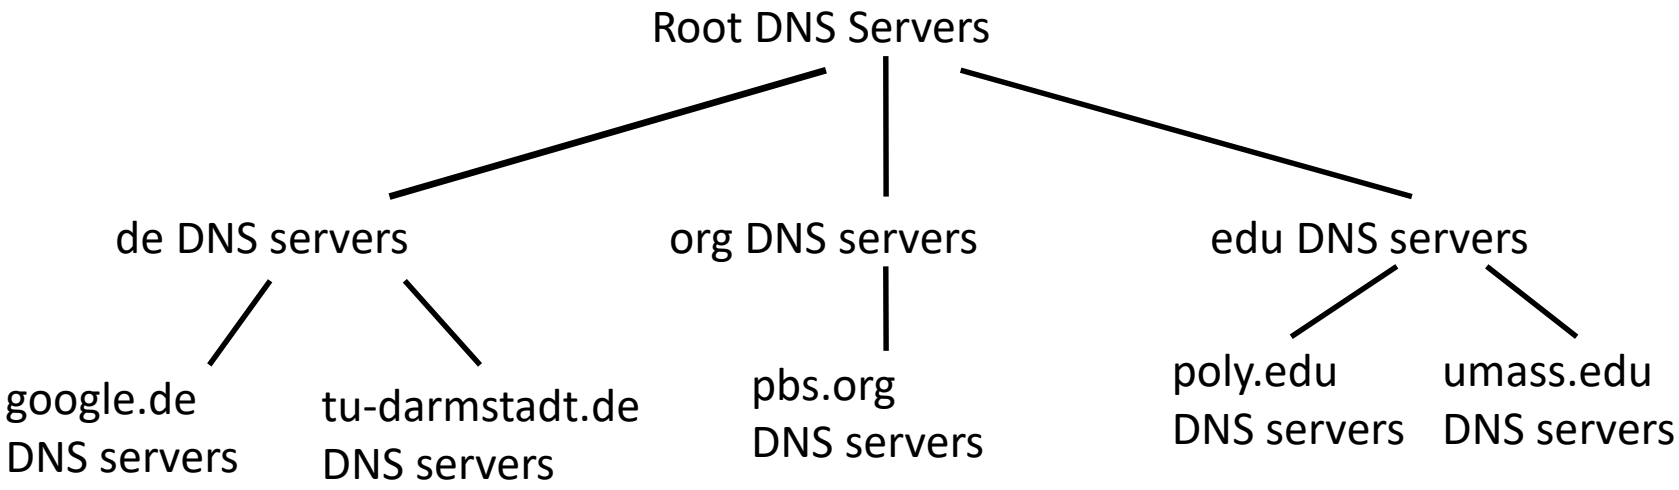

...and in reality...

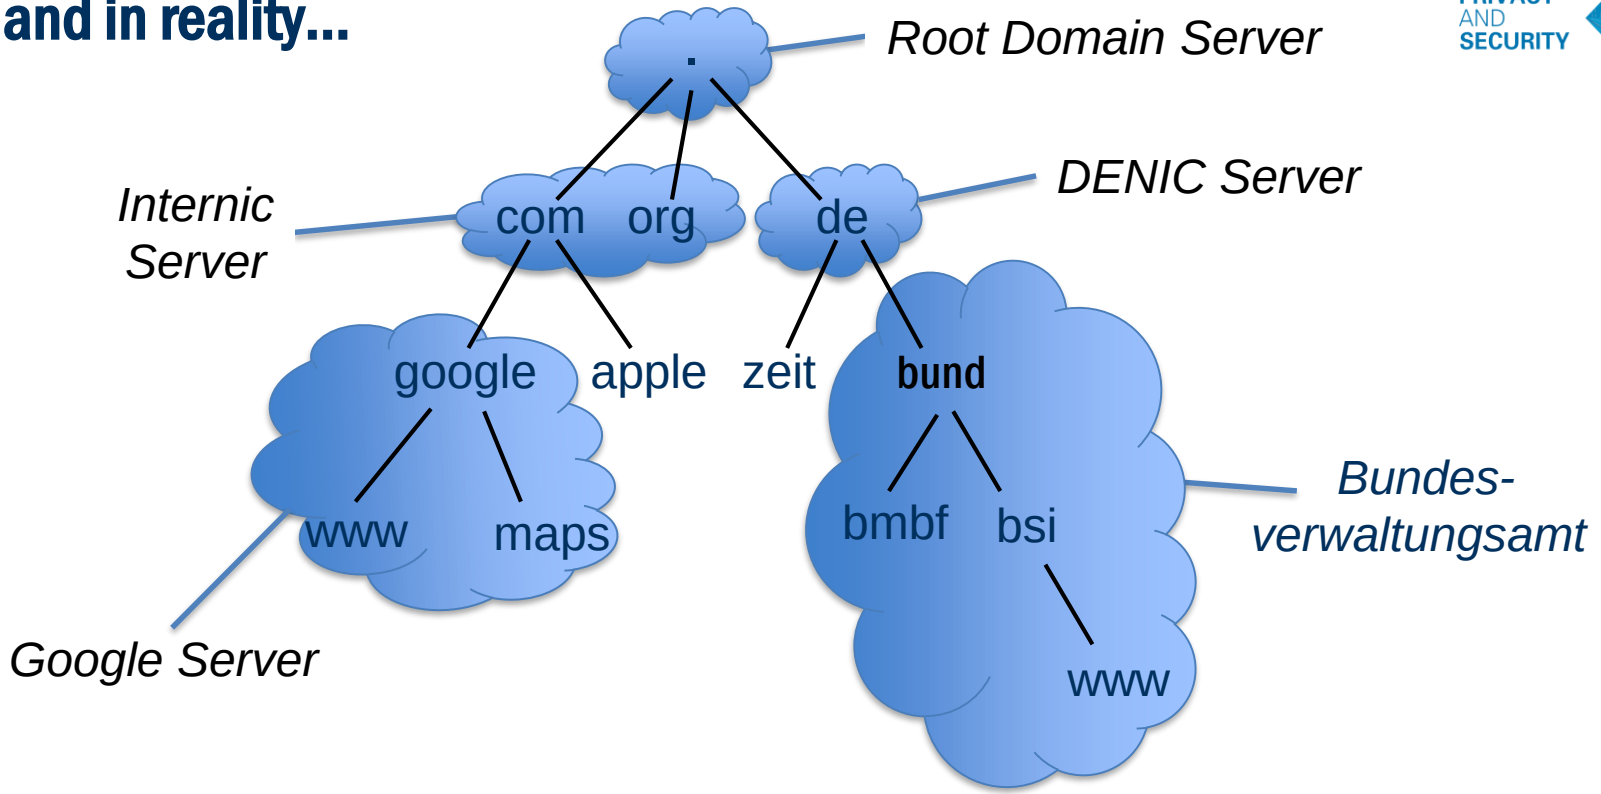

Client wants IP for www.dud.inf.tu-dresden.de; 1st approx:

- Client queries a root server to find **de** DNS server
- Client queries de DNS server to get **tu-dresden.de** DNS server
- Client queries tu-dresden.de DNS server to get IP address for www.dud.inf.tu-dresden.de

DNS: Root Name Servers

- Contacted by local name server that can not resolve name
- Root name server:
  - Contacts authoritative name server if name mapping not known
  - Gets mapping
  - Returns mapping to local name server

So, how many root nameservers are there actually? (physically)

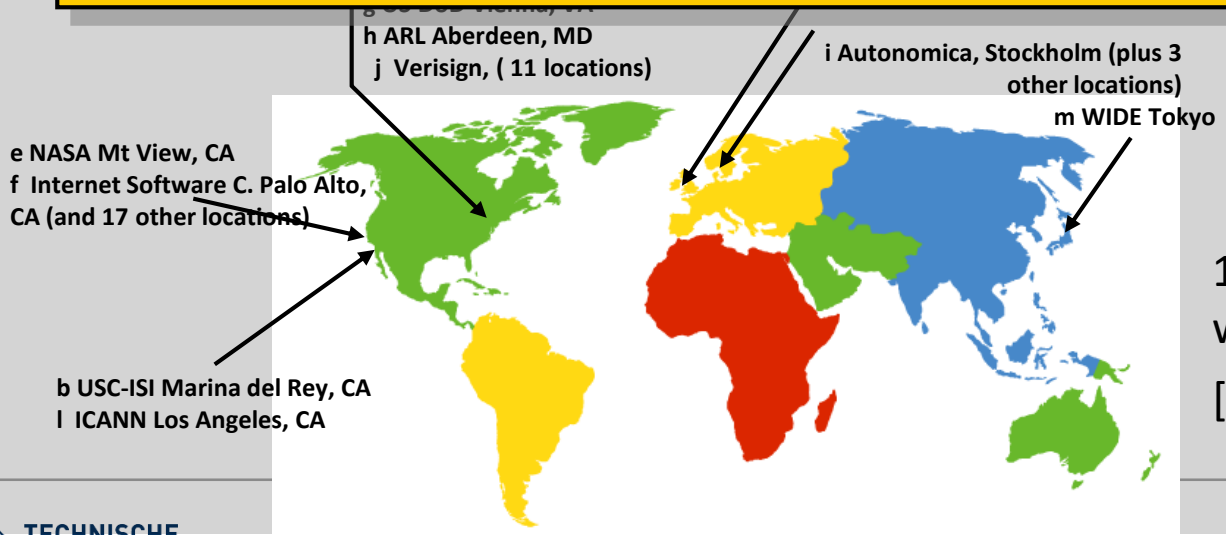

13 root name servers worldwide

[A..M].ROOT-SERVERS.NET

DNS: Root Name Servers

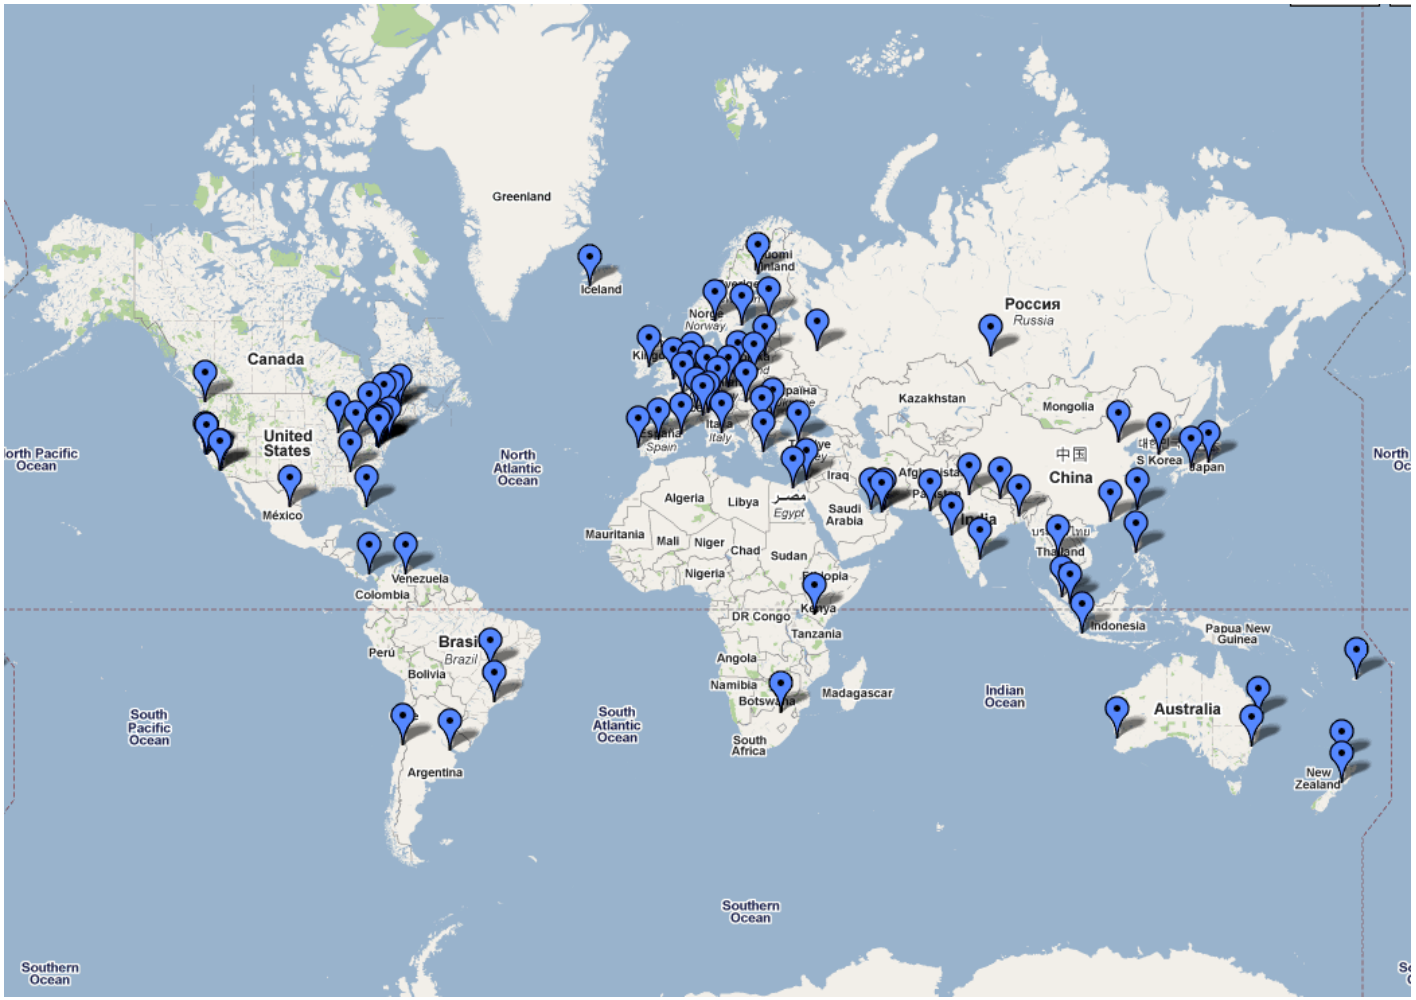

http://www.root-servers.org/

DNS – Components

Authoritative Server
- Server maintaining authoritative content of a complete DNS zone
- Top-Level-Domain (TLD) servers & auth servers of organization's domains
- Pointed to in parent zone as authoritative
- Possible load balancing: master/slaves

Recursive (Caching) Server
- Local proxy for DNS requests
- Caches content for specified period of time (soft-state with TTL)
- If data not available in the cache, request is processed recursively

Resolver
- Software on client's machines (part of the OS)
- Windows-* and *nix: Stub resolvers

- Delegate request to local server
- Recursive requests only, no support for iterative requests

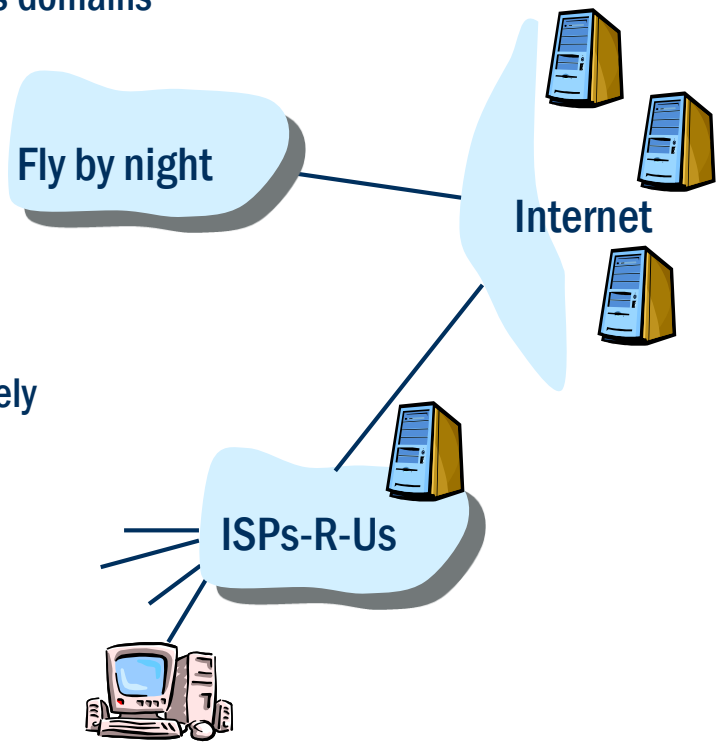

DNS – Resource Record Type

PRIVACY AND SECURITY

Atomic entries in DNS are called "Resource Records" (RR)

Format: <name> [<ttl>] [<class>] <type> <rdata>

- `name` (domain name of resource)
- `ttl` (Time-to-live)
- `class` (used protocol): IN (Internet), CH (Chaosnet)…
- `type` (record type): A (Host-Address), NS (Name Server),

  MX (Mail Exchange), CNAME (Canonical Name),

  AAAA (IPv6-Host-Address), DNAME (CNAME, IPv6)

- `rdata` (resource data): ***Content!*** *(What did we want to look up?)*

Specific DNS Records

RR Format: **(name, ttl, class, type, value)**

- Type=A
  - **`name`** is hostname
  - **`value`** is IP address
- Type=NS
  - **`name`** is domain (e.g. foo.com)
  - **`value`** is IP address of authoritative name server for this domain
- Type=MX
  - **`value`** is name of mailserver associated with **`name`**
- Type=CNAME
  - **`name`** is alias name for some "canonical" (the real) name

    `www.ibm.com` is really

    `servereast.backup2.ibm.com`
  - **`value`** is canonical name

DNS – Message Format

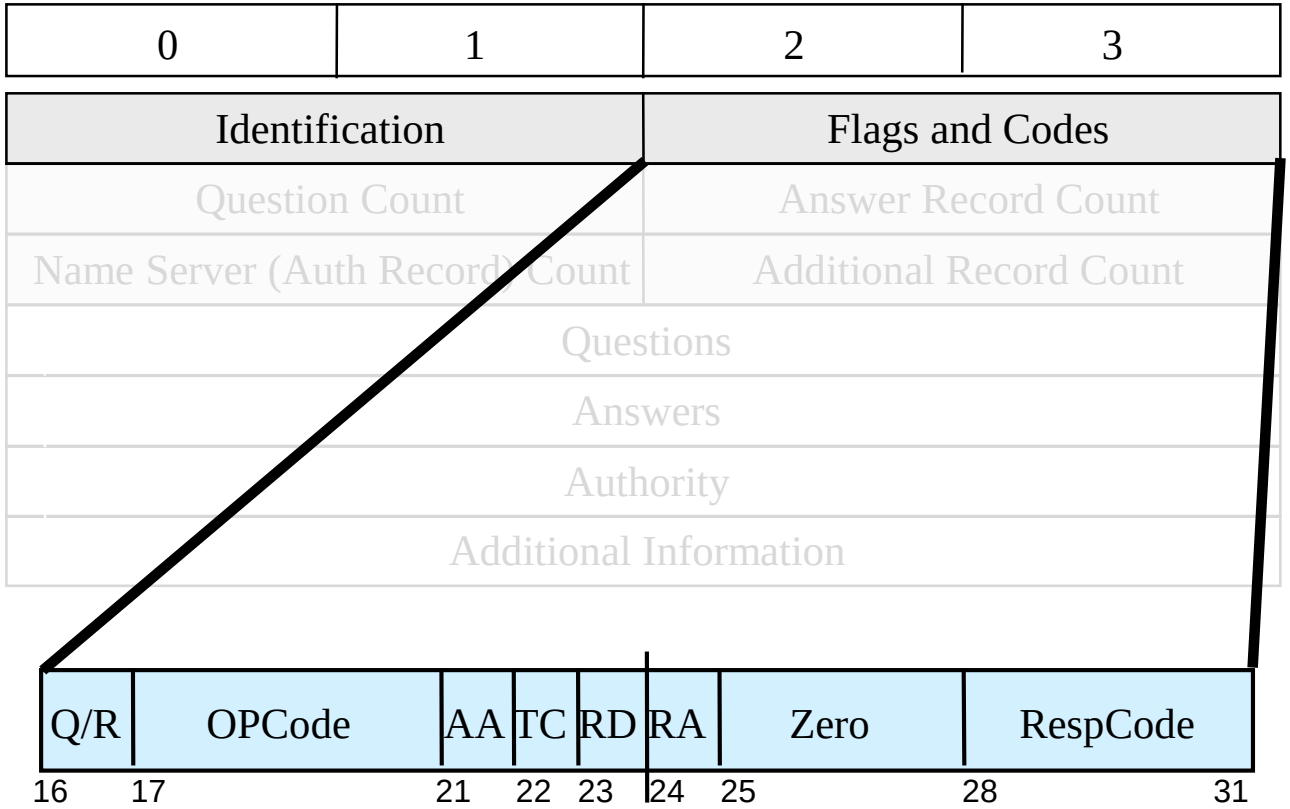

- Q/R *Query/Response Flag*
- *Operation Code*
- AA *Auth. Answer Flag*
- TC *Truncation Flag*
- RD *Recursion Desired Flag*
- RA *Recursion Available Flag*
- Zero (three resv. bits)
- *Response Code*

DNS – Header Fields

Identifier: a 16-bit identification field generated by the device that creates the DNS query. It is copied by the server into the response, so it can be used by that device to match that query to the corresponding reply

Query/Response Flag: differentiates between queries and responses (0 ~ Query, 1 ~ Response)

Operation Code: specifies the type of message (Query, Status, Notify, Update)

Authoritative Answer Flag (AA): set to1 if the answer is authoritative

Truncation Flag: When set to 1, indicates that the message was truncated due to its length (might happen with UDP, requestor can then decide to ask again with TCP as transport service)

Recursion Desired: set to 1 if requester desired recursive processing

Recursion Available: set to 1 if server supports recursive queries

DNS: Caching and Updating Records

Once (any) name server learns mapping, it caches mapping

- Stored as “soft state”: Cache entries timeout (disappear) after some time
- TLD servers typically cached in local name servers
- Thus, root name servers not often visited

Updating records, independent of TTL

- RFC 2136 defines dynamic updates
- BIND (>8) implements `nsupdate` (upon TSIG, see below)

Inserting Records Into DNS

- Example: just created startup “Fireblog”
- Register name `fireblog.de` at a registrar (e.g., denic)
  - Need to provide registrar with names and IP addresses of your authoritative name server (***primary*** and ***secondary***)
  - Registrar inserts two RRs into the de TLD server:

```
(fireblog.de, dns1.fireblog.de, NS)
(dns1.fireblog.de, 212.212.212.1, A)
```

- Add authoritative server Type A record for `www.fireblog.de` and Type MX record for `fireblog.de`

DNS – Recursive and Iterative Queries

DNS HEADER (send)
- Identifier: 0x3116
- Flags: 0x00 (Q)
- Opcode: 0 (Standard query)
- Return code: 0 (No error)
- Number questions: 1
- Number answer RR: 0
- Number authority RR: 0
- Number additional RR: 0

QUESTIONS (send)
- Queryname: (3)www(3)p2p(12)tu-darmstadt(2)de
- Type: 1 (A)
- Class: 1 (Internet)

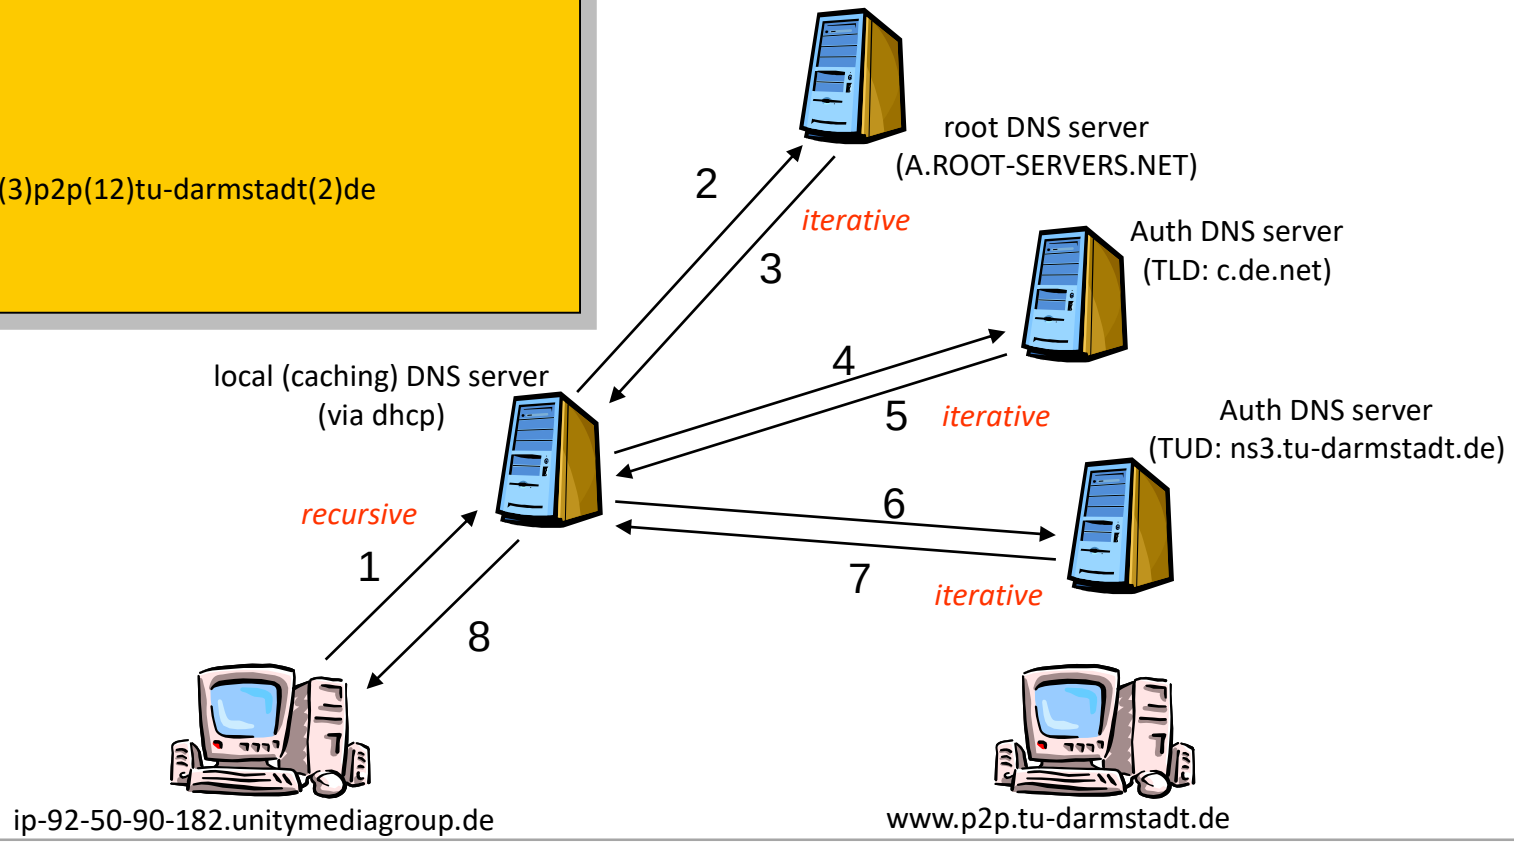

A Quick Example…

PRIVACY AND SECURITY

```
strufe@eris:~$ dnstracer -v www.p2p.tu-darmstadt.de
Tracing to informatik.tu-darmstadt.de[a] via 130.83.163.141, maximum of 3 retries
130.83.163.141 (130.83.163.141) IP HEADER
-Destination address:  130.83.163.141
-DNS HEADER (send)
-- Identifier:         0x3116
-- Flags:              0x00 (Q )
-- Opcode:             0 (Standard query)
-- Return code:        0 (No error)
-- Number questions:   1
-- Number answer RR:   0
-- Number authority RR:  0
-- Number additional RR: 0
-QUESTIONS (send)
-- Queryname:          (3)www(3)p2p(12)tu-darmstad
-- Type:               1 (A)
-- Class:              1 (Internet)
-DNS HEADER (recv)
-- Identifier:         0x3116
-- Flags:              0x8080 (R RA )
-- Opcode:             0 (Standard query)
-- Return code:        0 (No error)
-- Number questions:   1
-- Number answer RR:   2
-- Number authority RR:  0
-- Number additional RR: 0
-.......
```

```
QUESTIONS (recv)
- Queryname:          (3)www(3)p2p(12)tu-darmstadt(2)de
- Type:               1 (A)
- Class:              1 (Internet)
ANSWER RR
- Domainname:         (6)charon(7)dekanat(10)informatik(12)tu-darmstadt(2)de
- Type:               1 (A)
- Class:              1 (Internet)
- TTL:                1592 (26m32s)
- Resource length:    4
- Resource data:      130.83.162.6
ANSWER RR
- Domainname:         (3)www(3)p2p(12)tu-darmstadt(2)de
- Type:               5 (CNAME)
- Class:              1 (Internet)
- TTL:                49817 (13h50m17s)
- Resource length:    28
- Resource data:      (6)charon(7)dekanat(10)informatik(12)tu-darmstadt(2)de
Got answer [received type is cname]
```


DRESDEN concept

So where is the Info?

PRIVACY AND SECURITY

```
strufe@eris:~$ dnstracer -v -qns tu-darmstadt.de
Tracing to tu-darmstadt.de[ns] via 130.83.163.130
130.83.163.130 (130.83.163.130) IP HEADER
- Destination address:  130.83.163.130
DNS HEADER (send)
- Identifier:          0x4C45
- Flags:               0x00 (Q )
- Opcode:              0 (Standard query)
- Return code:         0 (No error)
- Number questions:    1
- Number answer RR:    0
- Number authority RR:  0
- Number additional RR: 0
QUESTIONS (send)
- Queryname:           (12)tu-darmstadt(2)de
- Type:                2 (NS)
- Class:               1 (Internet)
DNS HEADER (recv)
- Identifier:          0x4C45
- Flags:               0x8080 (R RA )
- Opcode:              0 (Standard query)
- Return code:         0 (No error)
- Number questions:    1
- Number answer RR:    5
- Number authority RR:  0
- Number additional RR: 9

.....
```

```
QUESTIONS (recv)
- Queryname:           (12)tu-darmstadt(2)de
- Type:                2 (NS)
- Class:               1 (Internet)
ANSWER RR
- Domainname:          (12)tu-darmstadt(2)de
- Type:                2 (NS)
- Class:               1 (Internet)
- TTL:                 70523 (19h35m23s)
- Resource length:     6
- Resource data:       (3)ns1(3)hrz(12)tu-darmstadt(2)de
ANSWER RR
- Domainname:          (12)tu-darmstadt(2)de
- Type:                2 (NS)
- Class:               1 (Internet)
- TTL:                 70523 (19h35m23s)
- Resource length:     5
- Resource data:       (2)ns(6)man-da(2)de
ANSWER RR
- Domainname:          (12)tu-darmstadt(2)de
- Type:                2 (NS)
- Class:               1 (Internet)
- TTL:                 70523 (19h35m23s)
- Resource length:     6
- Resource data:       (3)ns2(3)hrz(12)tu-darmstadt(2)de

.....
```

Answer ctd…

PRIVACY AND SECURITY

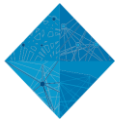

…….
ADDITIONAL RR
- Domainname: (3)ns1(3)hrz(12)tu-darmstadt(2)de
- Type: 1 (A)
- Class: 1 (Internet)
- TTL: 17335 (4h48m55s)
- Resource length: 4
- Resource data: 130.83.22.63

ADDITIONAL RR
- Domainname: (2)ns(6)man-da(2)de
- Type: 28 (unknown)
- Class: 1 (Internet)
- TTL: 38386 (10h39m46s)
- Resource length: 16
- Resource data: 2001:41b8:0000:0001:0000:0000:0000:0053

ADDITIONAL RR
- Domainname: (2)ns(6)man-da(2)de
- Type: 1 (A)
- Class: 1 (Internet)
- TTL: 38386 (10h39m46s)
- Resource length: 4
- Resource data: 82.195.66.249

ADDITIONAL RR
- Domainname: (3)ns2(3)hrz(12)tu-darmstadt(2)de
- Type: 28 (unknown)
- Class: 1 (Internet)
- TTL: 17335 (4h48m55s)
- Resource length: 16
- Resource data: 2001:41b8:083f:0022:0000:0000:0000:0063

…….

…….
ADDITIONAL RR
- Domainname: (3)ns2(3)hrz(12)tu-darmstadt(2)de
- Type: 1 (A)
- Class: 1 (Internet)
- TTL: 17335 (4h48m55s)
- Resource length: 4
- Resource data: 130.83.22.60

ADDITIONAL RR
- Domainname: (3)ns2(6)man-da(2)de
- Type: 1 (A)
- Class: 1 (Internet)
- TTL: 38386 (10h39m46s)
- Resource length: 4
- Resource data: 217.198.242.225

ADDITIONAL RR
- Domainname: (3)ns3(3)hrz(12)tu-darmstadt(2)de
- Type: 28 (unknown)
- Class: 1 (Internet)
- TTL: 17335 (4h48m55s)
- Resource length: 16
- Resource data: 2001:41b8:083f:0056:0000:0000:0000:0060

ADDITIONAL RR
- Domainname: (3)ns3(3)hrz(12)tu-darmstadt(2)de
- Type: 1 (A)
- Class: 1 (Internet)
- TTL: 17335 (4h48m55s)
- Resource length: 4
- Resource data: 130.83.56.60

Got answer

DNS – Lessons Learned

1. Structure name space (divide et impera)
2. Simple „routing“ b/c of structured (hierarchical) namespace
3. Store information at multiple locations
4. Maintain multiple connections
5. Be redundant! (Replicate…)
   - primary and secondary server, multiple TLD servers
6. Delegation using iterative or recursive forwarding

   (Btw: what are the pros and cons of each?)