

Chapter 2:

The Domain Name System

References:

- RRZN Hannover: Internet. Ein Einführung in die Nutzung der Internet-Dienste. Es gibt inzwischen die 8. Auflage. Erhältlich bei Beratung des URZ.
- D. Comer: Internetworking with TCP/IP. Prentice Hall, 1988, ISBN 0-13-470188-7.
- W. R. Stevens: TCP Illustrated, Vol. 1. Addison-Wesley, 1994, ISBN 0-201-63346-9.
- W. Richard Stevens: UNIX Network Programming, Vol. 1, 2nd Ed. Prentice Hall, 1998.
- Craig Zacker: Upgrading and Troubleshooting Networks — The Complete Reference. Osborne/McGraw-Hill, 2000, ISBN 0-07-212256-0, 918 pages.
- P. Mockapetris: Domain Names — Concepts and Facilities. RFC 1034, Nov. 1987.
- P. Mockapetris: Domain Names — Implementation and Specification. RFC 1035, 1987.
- E. Gavron: A Security Problem and Proposed Correction With Widely Deployed DNS Software. RFC 1535, October 1993, 5 pages.
- R. Elz, R. Bush: Clarifications to the DNS Specification. RFC 2181, July 1997, 15 pages.
- Florian Huber: Die sieben goldenen Domain-Regeln. <http://www.domain-recht.de>
- Uniform Domain Name Dispute Resolution Policy. <http://www.icann.org/udrp/udrp-rules-24oct99.htm>
- Holger Bleich: Ihre Suite im Internet, Hosting-Angebote im Vergleichstest. c't 10/2002, 112–121.

Objectives

After completing this chapter, you should be able to:

- explain how symbolic internet addresses (domain names) are structured.
- enumerate some types of records in the DNS.
- explain how the DNS works as distributed database.
- evaluate different domain/web hosting offers.

Overview

1. Name Space, Record Types
2. How the DNS works
3. The `nslookup` and `dig` Commands
4. Application Program Interface
5. Internet Hosting, Domain Laws

Foundations (1)

- The real IP-addresses are 32-bit numbers. But for humans, it is easier to work with/remember names.
- The Domain Name System (DNS) is a distributed database that manages the mapping from names to numbers (plus other information).
- Most internet applications use the DNS to translate host names into the IP-addresses.


E.g. Netscape prints “Looking up host: xyz” while the DNS-query runs. If it then prints “Netscape is unable to locate the server xyz”, the name xyz is not stored in the DNS (probably a typing error) or no DNS server can be reached (configuration error, network problem).

Foundations (2)

- At the beginning, host names were character strings without further structure (“flat namespace”).
- A file with the “Name → Number“ map was managed at InterNIC, every computer fetched a copy.
- But this works only with a small number of hosts:
 - ◇ The administration became too complicated because of the more and more frequent changes.
 - ◇ Name conflicts occurred more and more often.
 - ◇ The network load for the distribution of the file grows quadratically with the number of hosts.

Foundations (3)

- Solution: Hierarchical namespace, distributed DB.



Foundations (4)

- Similar tree structure as e.g. the UNIX file system.
- In contrast to file names, host names are written from the leaf towards the root (separated by periods), e.g. `“haendel.informatik.uni-halle.de.”`.
- A fully qualified/absolute domain name ends in a period (the name of the root node is empty).

The completion of other names depends on the DNS-Software. The usual solution is to append the local domain if the name does not contain a period, and otherwise to assume that the name is complete. Therefore, it is usually no problem if the period at the end is missing (depends on DNS software — I used a version of `nslookup` that always appended the local domain if one left out the `’.’` at the end).

Foundations (5)

- If one is on `“anubis.informatik.uni-halle.de”`, it normally does not work to specify `“www.mathematik”` to reach `“www.mathematik.uni-halle.de”`.
- Again, this depends on the DNS software.

Older software tried to attach portions of the current domain. Suppose one is currently on `a.b.com` and wants to reach `www.pitt.edu`. The software tried `“www.pitt.edu.a.b.com.”`, `“www.pitt.edu.b.com.”`, `“www.pitt.edu.com.”`, and `“www.pitt.edu.”` in this sequence. However, a domain `edu.com` was registered, and it would have been possible to redirect all traffic from `com`-hosts to `edu`-hosts to this domain.
- The “resolver” module (DNS query interface) is linked to the application software.

Foundations (6)

- Each name of a node in the tree is called a “label” .
- Such labels can be at most 63 characters long.

They should start with a letter and otherwise contain only digits and the hyphen “-”. The hyphen should not be the last character (see RFC 952 and RFC 1123). For instance, an underscore “_” is explicitly excluded. However, the DNS software should be able to work with arbitrary (non-empty) binary strings as labels (see RFC 2181). It is being discussed to permit national characters in domain names.

- The whole domain name may not be longer than 255 characters.
- The domain name comparison is not case-sensitive.

Foundations (7)

- “Top-Level Domains” (TLDs) are children of the root node, e.g.
 - ◇ `com, edu, gov, int, mil, nato, net, org` (“generic domains”)

The domains `com`, `org` and `net` can also be used by non-US organizations (at first, they were intended only for US organizations).
 - ◇ `de, at, ch, uk, fr, ca, jp` (“geographic domains”).

The geographic domains are the two character country codes defined in ISO 3166.
 - ◇ `arpa` (for inverse map, see below).

Foundations (8)

- New generic domains were selected by the ICANN board (see below) in November 2000:
`aero`, `biz`, `coop`, `info`, `museum`, `name`.
- Some country codes are also used by companies not really located in the country, e.g. `.to` (Tongo), `.tm` (Turkmenistan), `.tv` (Tuvalu), `.ag` (Antigua).
- Normally the leaf nodes are names of computers whereas inner nodes are proper domain names.

The DNS does not enforce this and there are many exceptions. In the DNS, every name that is generated by composing the labels from a node towards the root is called a domain name.

Foundations (9)

- Every node can contain different types of information (“resource records”).
- About 20 different record types were defined. The most important are:
 - ◇ **A** (Address): IP-Number.
 - ◇ **PTR** (Pointer): Name of the computer.

These records are used to map IP-numbers to host names (see below).
 - ◇ **CNAME** (Canonical Name): Real host name.

Used for aliases. E.g. under `www.informatik.uni-giessen.de` a **CNAME** record exists that refers to `odin.informatik.uni-giessen.de`.

Foundations (10)

- Types of DNS resource records, continued:
 - ◇ **HINFO** (Host Information): Two strings that describe hardware and operating system.

They are optional, and often not defined.
 - ◇ **MX** (Mail Exchanger): Host that acts as mail server for this domain (or this computer).

One can specify several computers with different priorities. Smaller numeric values mean that the computer is tried first.
 - ◇ **NS** (Name Server): Computer that manages the DNS records for this domain.

Foundations (11)

- Types of DNS resource records, continued:
 - ◇ **SOA** (Start of Authority): Administrative information for delegated domain (see below).

Contains the name of the primary name server, the email address of the name server administrator (with “.” instead of “@”), a serial number of the domain data, the time interval after which secondary name servers ask the primary name server for updates (REFRESH), the retry time interval (if the primary name server cannot be reached), how long secondary name servers may declare the data as authoritative if the refresh fails (EXPIRE), and the minimum time-to-live (how long data may be cached).

- ◇ **AAAA**: IPv6 address.

Foundations (12)

- The mapping of IP numbers to host names is done via a special domain `in-addr.arpa`.

- E.g. in the node

`50.14.48.141.in-addr.arpa`

a PTR-record with the following contents is stored:

`haendel.informatik.Uni-Halle.DE`

The IP-address of this machine is `141.48.14.50`. The DNS always lists labels from the local to the global, i.e. in the opposite sequence.

- Not all computers that are connected to the internet have DNS entries.

Overview

1. Name Space, Record Types
2. How the DNS works
3. The `nslookup` and `dig` Commands
4. Application Program Interface
5. Internet Hosting, Domain Laws

DNS as Distributed DB (1)

- The DNS is a distributed database: No computer (DNS server) contains the entire map.
- Distributed administration (delegation of subtrees).
- The DNS-tree is divided into “zones”: Parts that are administered at one place.

A zone is a subtree of the DNS tree (without subtrees that are zones by themselves). Root nodes of zones in the DNS tree are marked by records of the type **SOA** (“Start of Authority”). The zone mechanism can be compared with disks mounted into a UNIX file system.

- E.g. our computing center administers the zones **uni-halle.de** and **48.141.in-addr.arpa**.

DNS as Distributed DB (2)

- For each zone, there are at least two name servers (in order to protect against failures).

E.g. the root zone has 13 name servers, .de has 7.

- One of these name servers is the primary name server for the zone. At the primary name server, files with the data of the zone are administered.
- The other name servers are called “secondary name servers” of the zone. They contact e.g. every 3 or 6 hours the primary name server and copy its data.
- A server can be responsible for multiple zones.

DNS as Distributed DB (3)

- The domain `uni-halle.de` has four name servers:
`ns1.uni-halle.de`, `ns2.uni-halle.de`,
`ns3.uni-halle.de`, `deneb.dfn.de`.
- The name servers for the parent zone `de` contain `NS` (“name server”) records for `uni-halle.de` which refer to these four servers.
- The servers of the zone `uni-halle.de` contain an `SOA` record and further information about the domain (e.g. `MX`), its subdomains, and hosts.

Domain Registration (1)

- The name servers of the zone “de” are managed by the DENIC [<http://www.denic.de>].

Members of the DENIC association are internet service providers in Germany.

- If one is connected to the internet and has one's own name servers (or a contract with somebody who operates a name server) one can request the delegation of a domain below de to these servers.

The DENIC then stores NS-records for this domain in their name servers. This costs 116 Euro in the first year and 58 Euro in each following year. To their own members, DENIC offers special conditions.

Domain Registration (2)

- If one has no name server, one can request that DNS records (**A** and **MX**) are stored directly in the DENIC name servers.

The fee for up to 5 records is the same as that for the delegation of a domain. If one rents webspace from a DENIC member, one can get a domain much cheaper, see below.

- The DENIC also defines rules for acceptable domain names below **de**.

E.g. at least three characters that contain at least one letter, maximal length 63 characters, no names of top level domains, no city codes of car license numbers, no hyphen at the first, last, and third and fourth position.

Domain Registration (3)

- The ICANN (Internet Corporation for Assigned Names and Numbers) decides on top level domains [<http://www.icann.org/>].

For each top level domain, the ICANN accredits one (or more) registration companies.

- Domains below the generic top level domains were registered from 1993 to 1998 only by Network Solutions, Inc. [<http://www.networksolutions.com>].

Formally, this was the InterNIC.

Domain Registration (4)

- Now there are many registrars for domains below the generic top level domains `com`, `org`, `net`, see:
 - ◇ [<http://www.icann.org/registrars/accredited-list.html>]
 - ◇ [<http://www.internic.net>]
- Normally there is only one registrar for a domain.

The DENIC is the only registrar for the domain `de`, but several companies are members in the DENIC.
- IP numbers and `inaddr.arpa`-domains are managed in Europe by the RIPE NCC, see Chapter 1.

Until 1996 the DENIC had also this task.

Domain Registration (5)

- Domain registrars like DENIC have to collect the following information:
 - ◇ at least two name servers,
 - ◇ the owner of the domain (e.g. company),
 - ◇ an administrative contact,
 - A person who may decide things about the domain.
 - ◇ the person responsible for the bill,
 - ◇ a technical contact,
 - ◇ the administrator for the zone (nameserver).

Domain Registration (6)

- The data a registrar must collect can be queried with the `whois` command or a Whois web interface.
- Every registrar has its own database, and one must select the right database, e.g.
 - ◇ [<http://www.denic.de/servlet/Whois>]
The DENIC `whois` DB contains information about `.de` domains.
E.g. `whois -h whois.denic.de <Domain>`.
 - ◇ [<http://www.internic.net/whois.html>]
The Internic database contains information about `.com`, `.net`, `.org`, `.edu`, but only a reference to the database of the registrar.
 - ◇ “Universal Whois”: [<http://www.uwhois.com>].

Query Execution (1)

- The module that executes DNS queries is called a “Resolver” .

Under UNIX, it is a library that is linked to programs that use the DNS.

- The resolver must know at least one DNS server, to which it can send the query.

Under UNIX `/etc/resolv.conf` contains addresses of DNS servers. However, there are also other sources for the mapping from names to IP-numbers, e.g. `/etc/hosts` and the NIS/NIS+ database. Under Solaris, an entry in `/etc/nsswitch.conf` determines, which sources are queried in which sequence.

Query Execution (2)

- Example for `/etc/resolv.conf`:

```
domain informatik.uni-halle.de
nameserver 141.48.3.3
nameserver 141.48.3.51
nameserver 192.76.176.9
```

- Up to three name servers can be configured.

If there should be no answer from the first, the second is tried, and so on. Of course, IP-numbers must be specified, not names. The name servers in the example are: `ns1.uni-halle.de`, `ns3.uni-halle.de`, `deneb.dfn.de`.

- The `domain` entry is used to translate local names like `“haendel”` in absolute names.

Query Execution (3)

- If one uses a modem connection via PPP, name servers are normally automatically assigned

Under Windows, name servers can also be explicitly specified in the dialog box “Properties→Networking→TCP/IP Settings”.

- Example: Suppose we need to know the IP-address of `paradox.sis.pitt.edu`.
- The resolver sends the query to the first configured name server (`141.48.3.3`, `ns1.uni-halle.de`) via UDP. The name server listens on port `53`.

UDP (user datagram protocol) has less overhead than TCP.

Query Execution (4)

- There are two kinds of queries:
 - ◇ **Recursive:** The name server is asked to resolve the query completely, and to contact other name servers itself if necessary.

This is the usual case for DNS-clients (resolvers).

- ◇ **Iterative:** If the name server does not know the answer to the query, it sends back a reference to a name server that should be asked next.

This is the usual case for queries between name servers.

Query Execution (5)

- The name server of our computing center does not know the answer to the query, but it knows the addresses of the root servers of the DNS.

See [<ftp://ftp.rs.internic.net/domain/named.root>].

Name	IP-Number	Original Name / Comments
A.ROOT-SERVERS.NET.	198.41.0.4	NS.INTERNIC.NET
B.ROOT-SERVERS.NET.	128.9.0.107	NS1.ISI.EDU
C.ROOT-SERVERS.NET.	192.33.4.12	C.PSI.NET
D.ROOT-SERVERS.NET.	128.8.10.90	TERP.UMD.EDU
E.ROOT-SERVERS.NET.	192.203.230.10	NS.NASA.GOV
F.ROOT-SERVERS.NET.	192.5.5.241	NS.ISC.ORG
G.ROOT-SERVERS.NET.	192.112.36.4	NS.NIC.DDN.MIL
H.ROOT-SERVERS.NET.	128.63.2.53	AOS.ARL.ARMY.MIL
I.ROOT-SERVERS.NET.	192.36.148.17	NIC.NORDU.NET
J.ROOT-SERVERS.NET.	198.41.0.10	temporarily housed at NSI (InterNIC)
K.ROOT-SERVERS.NET.	193.0.14.129	housed in LINX, operated by RIPE
L.ROOT-SERVERS.NET.	198.32.64.12	temporarily housed at ISI (IANA)
M.ROOT-SERVERS.NET.	202.12.27.33	housed in Japan, operated by WIDE

Query Execution (6)

- The name server of our computing center sends the query to a root server, e.g. `A.ROOT-SERVERS.NET`.
- This happens to be responsible not only for the root domain, but also for the top level domain `edu`.
- Thus, it knows the name servers for the domain `pitt.edu` (which is delegated, i.e. a different zone):
`ns0-qip.ns.pitt.edu`, `ns1-qip.ns.pitt.edu`, ...
- The root server now sends names and IP-numbers of these three name servers back to the name server of our computing center.

Query Execution (7)

- The name server of our computing center sends the query now to `ns0-qip.ns.pitt.edu`.

We assume here that it uses an iterative query. The root name servers normally do not respond to recursive queries, but other name servers often accept them.

- This is responsible for `pitt.edu`, but the domain `sis.pitt.edu` is again a different zone.
- Thus, `ns0-qip.ns.pitt.edu` now sends the name servers responsible for `sis.pitt.edu` back to our name server: `icarus.lis.pitt.edu`, `acheron.lis.pitt.edu`, `thing.cs.pitt.edu`.

Query Execution (8)

- Next, the name server of our computing center sends the query to `icarus.lis.pitt.edu`.
- From there it finally gets the answer: “The IP-address of `paradox.sis.pitt.edu` is `136.142.116.28`”.
- In total, four DNS servers worked on the query.
- In order to reduce the network traffic caused by the DNS, all DNS servers contain a buffer (cache) for “resource records” that they recently received.

Query Execution (9)

- The administrator of a zone can define how long resource records of this zone may be buffered.
- E.g. if the name server of our computing center is asked again for the address of `paradox.sis.pitt.edu`, it returns the answer directly from its cache.

The client is told that this is a “non-authoritative answer”, as well as the address of the name server that has the “authoritative answer”.

- It has happened that name servers returned besides the requested data also falsified “resource records” that were also buffered and used for further queries.

Load Balancing via DNS

- If a service on the web is so popular that a single machine running a web server would be too slow, one can use several hosts with the same name.
- It is possible that a node in the DNS tree contains several resource records of the type “A”.

For gateways (connected to different networks) this is always the case.

- E.g. the DNS-server `ns1.altavista.com` returns 10 different addresses for `altavista.com`.

The addresses are returned each time in a different sequence. Most clients simply use the first address. In the way, the load is evenly distributed between the machines.

Server Configuration (1)

File "named.conf":

```
options {
    directory "/var/named";
};
zone "." {
    type hint;
    file "named.root";
};
zone "sis.pitt.edu" {
    type master;
    file "sis.zone";
};
zone "116.142.136.in-addr.arpa" {
    type master;
    file "sis.rev";
};
```

Server Configuration (2)

File "sis.zone":

```
@          SOA      icarus.sis.pitt.edu. mark.icarus.sis.pitt.edu. (
                2000102001 ; Serial YYYYMMDDNN
                10800 ; Refresh (3 hours)
                3600 ; Retry (1 hour)
                1209600 ; Expire (14 days)
                86400) ; Minimum TTL (1 day)

@          NS       icarus.sis.pitt.edu.
@          NS       acheron.sis.pitt.edu.
@          NS       thing.cs.pitt.edu.
@          MX 7     icarus.sis.pitt.edu.
@          MX 8     acheron.sis.pitt.edu.
www        CNAME    acheron.sis.pitt.edu.
icarus     A        136.142.116.2
acheron    A        136.142.116.10
paradox    A        136.142.116.28
```

Server Configuration (3)

File "sis.rev":

```
@ SOA icarus.sis.pitt.edu. mark.icarus.sis.pitt.edu. (
    2000102001 ; Serial YYYYMMDDNN
    10800 ; Refresh (3 hours)
    3600 ; Retry (1 hour)
    1209600 ; Expire (14 days)
    86400) ; Minimum (1 day)

@ NS icarus.sis.pitt.edu.
@ NS acheron.sis.pitt.edu.
@ NS thing.cs.pitt.edu.
2 PTR icarus.sis.pitt.edu.
10 PTR acheron.sis.pitt.edu.
28 PTR paradox.sis.pitt.edu.
```

Server Configuration (4)

File "named.root":

```
; formerly NS.INTERNIC.NET
.           3600000 IN NS A.ROOT-SERVERS.NET.
A.ROOT-SERVERS.NET. 3600000 A 198.41.0.4

; formerly NS1.ISI.EDU
.           3600000 NS B.ROOT-SERVERS.NET.
B.ROOT-SERVERS.NET. 3600000 A 128.9.0.107

; formerly C.PSI.NET
.           3600000 NS C.ROOT-SERVERS.NET.
C.ROOT-SERVERS.NET. 3600000 A 192.33.4.12
```

Overview

1. Name Space, Record Types
2. How the DNS works
3. The `nslookup` and `dig` Commands
4. Application Program Interface
5. Internet Hosting, Domain Laws

nslookup (1)

- Unter UNIX, the program `nslookup` can be used to query the DNS.

It also exists under Windows 2000, but not under Windows 98 or Windows ME. There is an interactive mode (if called without parameters) and a non-interactive mode (if the query is already specified on the command line). One can leave the program with `exit`. An overview of the commands gives `help`.

- `nslookup haendel.informatik.uni-halle.de`

```
Server: ns1.Uni-Halle.DE
```

```
Address: 141.48.3.3
```

```
Name: haendel.informatik.uni-halle.de
```

```
Address: 141.48.14.50
```

nslookup (2)

- `nslookup 141.48.14.50`

```
Server: ns1.Uni-Halle.DE
```

```
Address: 141.48.3.3
```

```
Name: haendel.informatik.Uni-Halle.DE
```

```
Address: 141.48.14.50
```

- `nslookup paradox.sis.pitt.edu`

```
Server: ns1.Uni-Halle.DE
```

```
Address: 141.48.3.3
```

```
Non-authoritative answer:
```

```
Name: paradox.sis.pitt.edu
```

```
Address: 136.142.116.28
```

nslookup (3)

- More information about non-authoritative answers:

```
nslookup -q=any paradox.sis.pitt.edu
```

```
Server: ns1.Uni-Halle.DE
```

```
Address: 141.48.3.3
```

```
Non-authoritative answer:
```

```
paradox.sis.pitt.edu internet address = 136.142.116.28
```

```
Authoritative answers can be found from:
```

```
sis.pitt.edu nameserver = acheron.sis.pitt.edu
```

```
sis.pitt.edu nameserver = icarus.sis.pitt.edu
```

```
sis.pitt.edu nameserver = thing.cs.pitt.edu
```

```
acheron.sis.pitt.edu internet address = 136.142.116.73
```

```
icarus.sis.pitt.edu internet address = 136.142.116.2
```

```
thing.cs.pitt.edu internet address = 136.142.80.5
```

nslookup (4)

- One can also select a specific name server, e.g.

```
nslookup paradox.... acheron.sis.pitt.edu
```

```
Server: acheron.sis.pitt.edu
```

```
Address: 136.142.116.73
```

```
Name: paradox.sis.pitt.edu
```

```
Address: 136.142.116.28
```

- In interactive mode, the server is selected with

```
NSLOOKUP> server acheron.sis.pitt.edu.
```

- Then one can enter queries to this server, e.g.

```
NSLOOKUP> paradox.sis.pitt.edu.
```

nslookup (5)

- One can specify which record types are requested:

```
nslookup -q=ns uni-halle.de.
```

```
Server: ns1.Uni-Halle.de
```

```
Address: 141.48.3.3
```

```
uni-halle.de nameserver = ns2.uni-halle.de
```

```
uni-halle.de nameserver = ns3.uni-halle.de
```

```
uni-halle.de nameserver = deneb.dfn.de
```

```
uni-halle.de nameserver = ns1.uni-halle.de
```

```
ns1.uni-halle.de internet address = 141.48.3.3
```

```
ns2.uni-halle.de internet address = 141.48.3.17
```

```
ns3.uni-halle.de internet address = 141.48.3.51
```

```
deneb.dfn.de internet address = 192.76.176.9
```

- In interactive mode, use “set querytype=ns”.

nslookup (6)

- With `-q=any` one gets in addition the `SOA`-record and the `MX`-records (“mail exchanger”).

```
...
uni-halle.de preference = 50, mail exchanger = mailgate.urz.uni-halle
uni-halle.de preference = 100, mail exchanger = mailgate2.urz.uni-hal
uni-halle.de
    origin = ns1.uni-halle.de
    mail addr = knauff.urz.uni-halle.de
    serial = 2004101401
    refresh = 10800 (3H)
    retry = 1800 (30M)
    expire = 604800 (1W)
    minimum ttl = 3600 (1H)
uni-halle.de nameserver = ns3.uni-halle.de
uni-halle.de nameserver = deneb.dfn.de
uni-halle.de nameserver = ns1.uni-halle.de
uni-halle.de nameserver = ns2.uni-halle.de
...
```

nslookup (7)

- Some name servers permit to list all entries in their domain:

```
UNIX> nslookup
> server regulus.informatik.uni-hannover.de.
> ls informatik.uni-hannover.de.
> exit
```

- Unfortunately, most administrators recently switched off this possibility.
- With the option `d2` query and answer are listed in complete detail.

Other DNS Lookup Tools (1)

- In the last time, `nslookup` has been criticized quite a lot for the following reasons:
 - ◇ `nslookup` does a reverse lookup for the nameserver to which it sends the real query (in order to print the name of the nameserver).

One can say that nobody asked `nslookup` to do this, and if it fails, it prevents the real query from being processed. If one uses `nslookup` to diagnose DNS problems, it is not unlikely that this will happen. Furthermore, `nslookup` asks the nameserver for its own name, but it might not know that, and if it is configured to answer no recursive queries, it cannot ask other servers.

Other DNS Lookup Tools (2)

- Problems of `nslookup`, continued:
 - ◇ `nslookup` might ask other name services, not only the DNS (e.g. `/etc/hosts`, NIS). It is not obvious where the answer came from.
 - ◇ `nslookup` has its own DNS client. Other programs are linked with a different DNS resolver library.

Therefore, other programs might still fail when `nslookup` works, or vice versa. E.g. the timeout/retry algorithm is different when there are several nameservers listed in `/etc/resolv.conf`.
 - ◇ It does not show all data received.

A chain of aliases that is interrupted somewhere is printed as “no answer”.

Other DNS Lookup Tools (3)

- Other tools for DNS lookup are `dig` and `host`.

They are available on our Linux and Solaris computers, but e.g. not under Windows XP.

- Example: `dig www.informatik.uni-halle.de`

- The output consists of several sections:

- ◇ First the version and the command are shown (can be switched on or off with `+[no]cmd`):

```
; <<>> DiG 9.3.5-P1 <<>> www.informatik.uni-halle.de
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 1379
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 4, ADDITIONAL: 4
```

Other DNS Lookup Tools (4)

- Sections of `dig` output, continued:
 - ◇ Next, a summary of the results is given (can be switched on or off with `+[no]comments`, this also influences section headers):

```
;; Got answer:  
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 1379  
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 4,  
                                         ADDITIONAL: 4
```

- ◇ Then the query is shown (`+[no]question`):

```
;; QUESTION SECTION:  
;www.informatik.uni-halle.de. IN A
```

Other DNS Lookup Tools (5)

- Sections of `dig` output, continued:
 - ◇ The answer follows (`+[no]answer`): this also influences section headers):

```
;; ANSWER SECTION:  
www.informatik.uni-halle.de. 3600 IN A 141.48.3.149
```

(3600 is the remaining time-to-live for this entry).

- ◇ Next, the DNS servers with the authoritative answer are listed (`+[no]authority`):

```
;; AUTHORITY SECTION:  
uni-halle.de. 3600 IN NS deneb.dfn.de.  
uni-halle.de. 3600 IN NS ns3.uni-halle.de.  
uni-halle.de. 3600 IN NS ns2.uni-halle.de.  
uni-halle.de. 3600 IN NS ns1.uni-halle.de.
```

Other DNS Lookup Tools (6)

- Sections of `dig` output, continued:
 - ◇ Additional DNS records are listed, typically the IP numbers of the DNS servers (`+[no]additional`):

```
;; ADDITIONAL SECTION:  
ns1.uni-halle.de. 14400 IN A 141.48.3.3  
ns2.uni-halle.de. 14400 IN A 141.48.3.17  
ns3.uni-halle.de. 14400 IN A 141.48.3.51  
deneb.dfn.de. 32374 IN A 192.76.176.9
```

- ◇ Finally, statistics are shown (`+[no]stats`):

```
;; Query time: 2 msec  
;; SERVER: 141.48.3.3#53(141.48.3.3)  
;; WHEN: Thu Apr 30 09:42:50 2009  
;; MSG SIZE rcvd: 203
```

Other DNS Lookup Tools (7)

- One can switch off all sections with `+noall` and then switch on selectively certain sections:

```
dig +noall +answer www.acm.org
```

- There is also an option `+short` for a short output.
- One can query for special resource record types:

```
dig +noall +answer MX informatik.uni-halle.de
```

- The resource record type `ANY` works also.
- `+trace` shows in detail how the query is resolved.
- `+search` adds the local domain name if necessary.

Overview

1. Name Space, Record Types
2. How the DNS works
3. The `nslookup` and `dig` Commands
4. Application Program Interface
5. Internet Hosting, Domain Laws

Program Interface (1)

- The function `gethostbyname` returns for a given name a pointer to a structure with information about a host:

```
struct hostent {
    char    *h_name;          /* official name of host */
    char    **h_aliases;     /* alias list */
    int     h_addrtype;      /* host address type */
    int     h_length;        /* length of address */
    char    **h_addr_list;   /* list of addresses */
};
#define h_addr h_addr_list[0] /* first address */
```


Program Interface (2)

- The structure is defined in `netdb.h` under UNIX and `winsock.h/winsock2.h` under Windows.

One must include this header file.

- Example:

```
struct hostent *h;  
h = gethostbyname("www.sis.pitt.edu");
```

- The function returns a null pointer in case of errors.

Under UNIX, a new global variable `int h_errno` is used instead of `errno` (which is used by other socket functions). It can have values like `HOST_NOT_FOUND` and `NO_DATA` (valid domain, but no A record) which are defined in `netdb.h`. New implementations have `hstrerror`. Under Windows, one can use `WSAGetLastError()` as for the other functions.

Program Interface (3)

- For IP-addresses, `h->h_addrtype` is `AF_INET`.

For safety, check also that `h->h_length == sizeof(struct in_addr)` before calling `memcpy` below. For IPv6, `h->h_addrtype` will be `AF_INET6` and `h->h_length` will be 16.

- Then, if `addr` is declared as `struct sockaddr_in`, one can copy the IP address with

```
memcpy(&(addr.sin_addr), h->h_addr_list[0],  
      h->h_length);
```

This works because `gethostbyname` returns numbers in the network byte order.

Program Interface (4)

- The function `gethostbyaddr` returns the same structure for a given IP-number:

```
struct hostent *h;
struct in_addr a;
a.s_addr = inet_addr("134.176.28.60");
if(a.a_addr == INADDR_NONE)
    /* Invalid address (or broadcast) */ ...
h = gethostbyaddr(&a, sizeof(a), AF_INET);
if(h == 0) /* Error */ ...
```

- Use `uname` or `gethostname` followed by `gethostbyname` to determine the local IP address.

Overview

1. Name Space, Record Types
2. How the DNS works
3. The `nslookup` and `dig` Commands
4. Application Program Interface
5. Internet Hosting, Domain Laws

Internet Hosting (1)

- If one is connected to the internet, one can operate one's own web server.

Of course, one should to be connected around the clock in order for the web server to be reachable without restrictions. A flatrate is not intended for this application, e.g. the Telekom automatically disconnects their flatrate customers after 24 h connection time. One can immediately reconnect, but is assigned a new IP address. There are dynamic DNS servers which permit to change the mapping from names to IP-addresses easily, so that one can keep a stable name.

- An alternative is to rent space on a web server (WWW Hosting).

Many online services and ISPs include some webspace, but only in a subdirectory of their web server, not under one's own domain.

Internet Hosting (2)

- It is possible to apply for a domain and let this domain name refer to a host on which one has rented web space.
- Often many domains will point to the same host.
 - For private applications, it would be too expensive to rent an entire computer.
- If the host has as many IP numbers, it can use the contacted IP address for the decision which web page should be delivered.
 - Although it is really only one machine with a single web server, it looks like many machines, each with a web server running.

Internet Hosting (3)

- But IP numbers are a scarce resource. Therefore, today many different domain names are mapped to the same IP number.
- In HTTP/1.1, the request for a web page contains the name of the server, and therefore a single machine can look like many different web servers, although all have the same IP number.

With HTTP/1.0, it does not work. There the web server does not get the information about the requested host name. The protocol designers assumed that the WWW-server knows its own host name.

Web-Hosting Offers (1)

- 1&1 WebHosting: [<http://www.puretec.de>]
- ◇ Web Address 4.0: 0.69 Euro/Month.

All 1&1 prices are plus a 9.60 Euro one-time fee. This package does not include webspace, but one becomes the domain owner. Only [.de](#) domains can be chosen in this package. One can redirect web page requests to a different URL. One email account is included. The package also includes Email-support, and also telephone support (0800) for the first four weeks.

Web-Hosting Offers (2)

- 1&1 WebHosting, continued:
 - ◇ Basic 4.0: 3.99 Euro/Month.

If one accepts advertisements on one's web pages, it costs only 2.99 Euro/month. This package includes: 1 domain (.de, .com, .net, .org), 25 MB Webspace, FTP access for uploading web files, 10 email accounts, 100 additional aliases, 4000 MB transfer volume/month (each additional MB costs 1.5 cent), Selection from given simple CGI programs, 1&1 WebChat, WAP support, possibility to redirect page requests to a different URL. The package includes NetObjects Fusion MX, PhotoImpact 7 SE, and other software (one has to pay 6 Euro for shipping and handling).

Web-Hosting Offers (3)

- 1&1 WebHosting, continued:

- ◆ Star Package 4.0: 6.99 Euro/Month.

Includes: 2 domains, 5 subdomains, domain redirection, 50 email accounts, 500 email aliases, 100 MB webspace, 7500 MB monthly transfer volume, access to apache logfiles, selection from given CGI programs (e.g. counter), access to a graphic library.

Web-Hosting Offers (4)

- 1&1 WebHosting, continued:

- ◇ Profi Package 4.0: 12.99 Euro/Month.

Three domains (including .info), 10 subdomains, 250 MB web-space, 15 GB transfer volume/month, Server Side Includes (SSI), PHP3, PHP4 (embedded into HTML web pages), MS Frontpage 2002 Extensions, 1&1 WebDatabase.

- ◇ Premium Package 4.0: 24.99 Euro/Month.

5 domains (including .info, .biz), 20 subdomains, 500 MB web-space, 25 GB monthly transfer volume, SSL support ([https:](https://)), arbitrary CGI programs (Perl, Python, PHP3, PHP4), MySQL database.

Web-Hosting Offers (5)

- 1&1 WebHosting, continued:
 - ◇ Power Package 4.0: 39.99 Euro/Month.

10 domains (including .info, .biz), 50 subdomains, 1 GB web-space, 50 GB monthly transfer volume, SSL, CGI, MySQL as before, CRON jobs, SSH access (login on the server).
 - ◇ Own Server: 69 Euro/Month.

2000 MHz Celeron, 256 MB RAM, 40 GB disk, 75 GB monthly transfer volume. A Linux system with root access costs only 49 Euro/Month. Both require a one-time fee of 99 Euro. Of course, in this way one gets an own IP address.

Web Hosting Offers (6)

- **Strato Medien AG:** [<http://www.strato.de>]
- ◇ **Web Visitenkarte A:** 0.28 Euro/Month

Plus 9.90 Euro one-time fee. 1 web page can be constructed with an online form (no FTP access), or the domain can be redirected to another web server. No POP3 mailbox, but one addresses can be forwarded to other accounts. Data transfer 4 GB, each additional MB costs 1.5 Cent.

Web Hosting Offers (7)

- Strato Medien AG, continued:
 - ◇ Web Visitenkarte S: 0.99 Euro/Month

Plus 19.90 Euro one-time fee. Includes one `.de`-domain, 2 MB webspace (daily backups), FTP access for uploading web files, domain cannot be redirected, 4 GB traffic, 12 POP3 mailboxes, 120 email aliases, telephone support costs 1.86 Euro/min, no email support, integrated messaging center (Fax/Email, 100 messages/month included), Namo WebEditor 4, Paint Shop Pro 4.15, PromoWare 2.1, GS ShopBuilder Entry, graphics archive.

Web Hosting Offers (8)

- **Strato Medien AG, continued:**

- ◇ **PowerWeb A: 4.99 Euro/Month.**

3 domains (.de, .com, .net, .org), 85 MB webspace, 85 POP3 mailboxes, 850 email aliases, 7.5 GB data transfer/month (each additional MB costs 0.015 Euro), server statistics, standard CGIs (about 20), GS ShopBuilder Entry 2, online shop CGI without limit on number of articles, WAP homepage, Adobe GoLive 6.0, Adobe Photoshop Elements 2.0, and other software.

- ◇ **PowerWeb M: 9.90 Euro/Month**

3 domains, 170 MB webspace, 170 POP3 mailboxes, **unlimited data transfer**, otherwise as before.

Web Hosting Offers (9)

- **Strato Medien AG, continued:**

- ◇ **Premium S: 19.90 Euro/Month.**

5 domains, 300 MB webspace (data mirrored on a RAID system, daily backups), 300 POP3 mailboxes, 3000 email aliases, 25 GB data transfer/month (0.04 Euro per additional MB), SSL encoding, own CGI programs, MS Frontpage 2002 Server Extensions, RealMedia Server, Live Picture Server (FlashPix), PHP 3, PHP 4, Python, SSI, MySQL database, telnet access, anonymous FTP (incoming FTP), telephone support for 0.12 Euro/Min, Adobe GoLive 6.0, Adobe Photoshop Elements 2.0, Adobe LiveMotion 2.0.

- ◇ **Additional .de domain: 0.99 Euro/Month.**

Additional .com, .org, .net domain: 1.45 Euro/Month.

.at/.ch domain: 3.99 Euro/Month plus 99 Euro one-time fee.

Web Hosting Offers (10)

- Loomes: [<http://www.loomes.de>]
- HostEurope: [<http://www.hosteurope.de>]
- Domainfactory: [<http://www.domainfactory.de>]
- Domainbox: [<http://www.domainbox.de>]
- FreeCity: [<http://domains.freecity.de>]

The domain registration costs a one-time fee of 14.99 Euro. If one accepts advertisements (via pop-up window) for all accesses to the web pages under the domain, there is no further fee.

- Redirect to you: [<http://www.r2u.de/>]

Selection Criteria (1)

- Will one be registered as domain owner and administrative contact for the domain at DENIC?

Important for changing to a different web hosting company later. Even if one is the domain owner, there may be a fee for the domain transfer.

- Can arbitrary web pages and other files be stored?

Some cheap offers have only one webpage with a fixed format. Otherwise, one gets FTP access to upload arbitrary files on the web server.

- How many domains? Also `.com`, `.net`, `.org`, etc?
Subdomains?

Selection Criteria (2)

- How much disk space for files on the web server (“webspaces”)?

- Data transfer volume per month?

If the pages are accessed very often (e.g. from a robot operated by some hacker), that can cause large extra costs. Some companies offer to simply switch the website off when the quota is reached, others only send a warning email, others maybe not even that.

- Does one get statistical data about web accesses?

Nice graphical representation? How detailed (e.g. server log files)?

- Redirection to an existing web server?

Selection Criteria (3)

- Can one use CGI programs?

Only select one from a fixed collection or write one's own programs? If yes, what languages are supported? Telnet/SSH access would be useful for debugging the programs, but is not strictly needed. The operating system might be interesting.

- Can the CGI programs access a database?

If yes, what DBMS? Are there restrictions in the SQL language (e.g. MySQL)? Does it have support for transactions? What about the safety of the data: Does the DBMS write log files, does anybody make backup copies?

- Online shop included?

Selection Criteria (4)

- SSL-support ([https:](https://))?

Passwords and credit card numbers should be encrypted while sent through the internet.

- Can certian pages be password-protected?

- Email Accounts?

POP3 mailboxes, email forwarding, autoresponder? How large can the emails and their attachments be? How large is the mailbox? Are emails automatically deleted after some time? Can an SMS be sent for incoming emails (and what does this cost)? Can emails be sent via fax?

Selection Criteria (5)

- WAP-support?
- Is software included?

Some software is shipped on a CD, other software has to be used via the web on the server of the web hosting company. Sometimes the software is much more expensive if bought separately.

- Technical support in case of problems?

Sometimes this is very expensive (e.g. telephone support via 0190-number). Is there a hotline 7 days a week, 24 h a day?

Selection Criteria (6)

- How available is the web server?

If the web server is often not reachable/down, this is a problem. E.g. are there redundant internet connections? Do they have battery backup for power failures?

- How fast is the internet connection of the server?

If most customers live in Germany, the server should probably be in Germany with a fast connection to the DE-CIX.

- Does the web hosting company make backups of the data on the web server? Do they use RAID-systems?

Domain Laws (1)

- See: [<http://www.domain-recht.de/>]
- When one registers a domain, one has to sign that one does not violate rights of other persons or companies, and that one will pay the fees for lawsuits.

This is probably the main reason why one is today normally registered as domain owner.

- Trademarks/company names may not be registered, even if by chance it is ones own family name.

Also combinations like microsoft-haters.de can be problematic as well as small changes like microsaft.de.

Domain Laws (2)

- Names of persons belong to these persons.
- Titles of journals, books, software, films are protected if they are very well known or there is a possibility of confusion.
- Names of cities or countries belong to them.
In addition, DENIC does not permit auto license city codes.
- Names of government agencies belong to them.
- “Typing error domains” like `aliavista.com` belong to the owner of the corresponding well known domain `altavista.com`.

Domain Laws (3)

- General descriptive names like `database-course.de` are not forbidden.

Such domains gets the first person who requests them, even though there are other database courses.

- However, `mitwohnzentrale.de` was successfully sued.

The appeal is still open and some experts feel that the judge made an error. This would be a problem for many domains, e.g. also `buch.de`.

- Domains are sold, sometimes for large sums.

The internet community thinks that this is an abuse. Judges assumed for quite some time that trading domains is immoral. The domain `loans.com` was sold for 3 Mio \$.

Domain Laws (4)

- The rules for domain names depend on the top-level domain (they are determined by the registrar).

At least, the registrar has the technical possibility to change the name server entries. Of course, one can also go before a usual court.

- Some domains were lost because bills of the registrar were not paid in time.
- All registrars for the domains `.com`, `.org`, `.net` have adopted the “Uniform Domain Name Dispute Resolution Policy”.

[\[http://www.icann.org/udrp/udrp-rules-24oct99.htm\]](http://www.icann.org/udrp/udrp-rules-24oct99.htm)

Domain Laws (5)

- The UDRP requires that a domain is transferred if the complainant proves three things:
 - ◇ The domain name is identical or confusingly similar to a trademark owned by the complainant.
 - ◇ The current owner of the domain has no rights or legitimate interests in the domain.
 - ◇ The domain was registered with evil purpose.

E.g. in order to sell it, to get page hits because of the confusion, or to disturb the business of the complainant.