ALBERT-LUDWIGS-
UNIVERSITÄT FREIBURG

Communication Systems

DNS

University of Freiburg
Computer Science
Computer Networks and Telematics
Prof. Christian Schindelhauer

CoNe Freiburg

IIF
INSTITUT FÜR INFORMATIK FREIBURG

What is DNS?

- Imagine: Try to remember the telephone numbers of your friends instead of their names
- What is DNS? - What Internet users use to reference anything by name on the Internet
- The mechanism by which Internet software translates names to addresses and vice versa
- A lookup mechanism for translating objects into other objects
- A globally distributed, loosely coherent, scalable, reliable, dynamic database

DNS – the Internet Telephony Book

- 1970's ARPANET
  - Host.txt maintained by the SRI-NIC
  - pulled from a single machine
  - Problems
    - traffic and load
    - Name collisions
    - Consistency
- DNS created in 1983 by Paul Mockapetris (RFCs 1034 and 1035)
- Modified, updated, and enhanced by a myriad of subsequent RFCs (e.g. 3490-2)

DNS – Features

- A lookup mechanism for translating objects into other objects
- A globally distributed, loosely coherent, scalable, reliable, dynamic database
- Comprised of three components
  - A “name space”
  - Servers making that name space available
  - Resolvers (clients) which query the servers about the name space
- Data is maintained locally, but retrievable globally
  - No single computer has all DNS data
- DNS lookups can be performed by any device and any service
- Remote DNS data is locally cachable to improve performance

DNS – as an IP Service

- DNS is an IP based service
  - the IP world can live without DNS (the humans may not), but the DNS is dependent of IP
- DNS is application level protocol like others, e.g. HTTP, SSH, DHCP, ...
- Mostly using UDP as transport layer protocol, maximum DNS UDP packet size is 512Byte (restricts the size of DNS replies)
  - too long answers are truncated (client is told by truncate flag)
- Uses well-known port 53 for client-server-interaction, see e.g. /etc/services in Unix-like systems for the list of ports

Loose Coherency

- The database is always internally consistent
  - each version of a subset of the database (a zone) has a serial number
  - serial number is incremented on each database change
- Changes to the master copy of the database are replicated according to timing set by the zone administrator
- Cached data expires according to timeout set by zone administrator

Scalability

- No limit to the size of the database
  - One server may have over 20,000,000 names
  - Not a particularly good idea
- “No limit” to the number of queries
  - 50,000 queries per second handled easily
- Queries distributed among masters, slaves, and caches
  - principles are explained little bit later

Reliability

- Data is replicated
  - Data from master server my be copied to several slaves
- Clients can query
  - master server
  - any of the copies at slave servers
  - use several caches
- Clients will typically query local caches first
  - see your DSL/cable router for DNS server assignments
  - e.g. local server for Freiburg university campus is 132.230.200.200 and 132.230.200.201 is caching server and server for uni-freiburg.de.
  - but you are free to contact e.g. the Freiburg university server

Dynamics

- Database can be updated dynamically
  - add/delete/modify of almost any record
  - example: www.dyndns.org and several other similar services use this characteristic
    - very short setting of TTL used
    - typically only one direction of name resolution – from name to IP
    - integrated in many IAD (Internet Access Devices – Telco lingo)
- Modification of the master database triggers replication
  - only master can be dynamically updated
  - thus creates a single point of failure

Concepts

- The name space needs to be made hierarchical to be able to scale
  - The idea is to name objects based on
  - location (within country, set of organizations, set of companies, etc)
  - unit within that location (institute within a faculty)

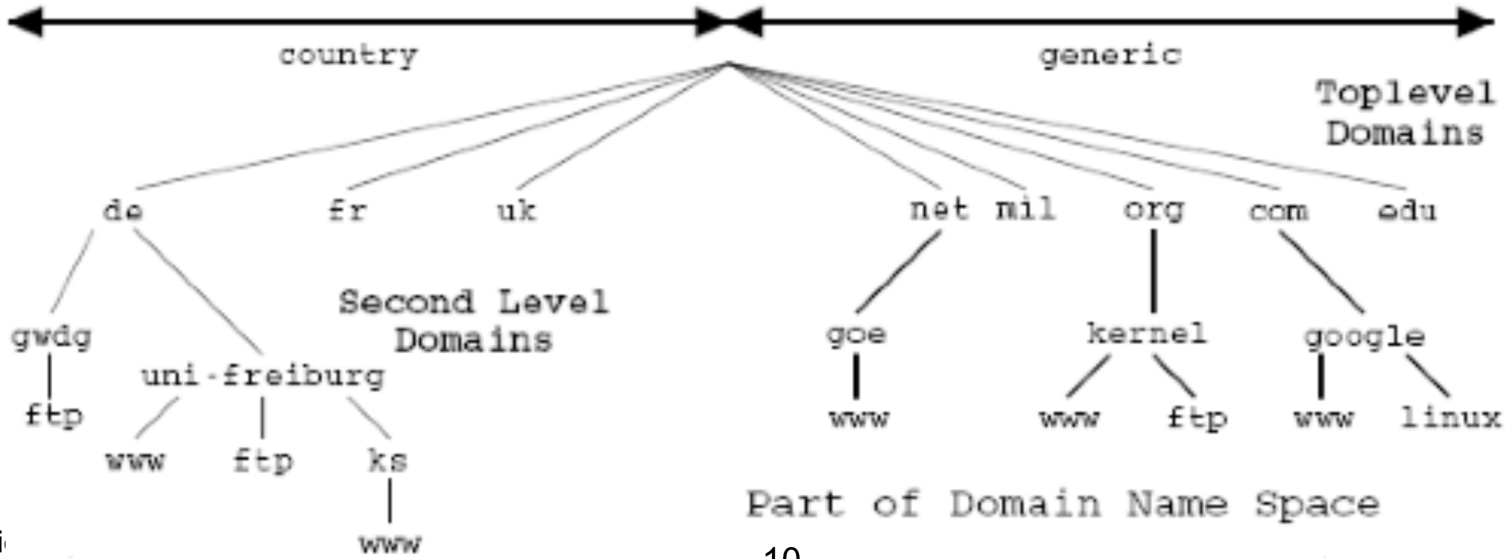

Naming within DNS

- Fully Qualified Domain Name (FQDN) of a specific host
- WWW.KS.UNI-FREIBURG.DE.
- Labels separated by dots
  - concept known from dotted quad notation of IP addresses (good readable representation of objects for humans)
  - given example not a host by definition. e.g.
    - www.rz.uni-freiburg.de (hostname – webserver within the “subdomain” of the Comp. Dept.)
    - rz.uni-freiburg.de (hostname – mailserver for the Comp. Dept. but subdomain name in the same moment)
- DNS provides a mapping from FQDNs to resources of several types
- Names are used as a key when fetching data in the DNS

Naming System and Conventions

- Domain names can be mapped to a tree
- New branches at the ‘dots’
- No (real) restriction to the amount of branches
  - www.ks.uni-freiburg.de
  - ftp.uni-freiburg.de
  - www.google.de
  - electures.informatik.uni-freiburg.de
- Domains are “namespaces”
  - Everything below .de is in the de domain
  - Everything below uni-freiburg.de is in the uni-freiburg.de domain and in the de domain

Concepts - Namespace

- Each node has a label
  - The root node has a null label, written as "."
- Each node in the tree must have a label
  - A string of up to 63 (8 bit) bytes
- The DNS protocol makes NO limitation on what binary values are used in labels
  - RFCs 852 and 1123 define legal characters for "hostnames"
    - A-Z, 0-9, and "-" only with a-z and A-Z treated as the same
    - internationalization (IDNA: "umlaut", chinese character, ... domains) were defined in 2003 (RFC 3490)
    - int. names are made compatible (normalized) via nameprep algorithm (RFC 3491) and then via punycode (RFC 3492) translated to the allowed DNS character set

Concepts – Domain Name

- Sibling nodes must have unique labels
- The null label is reserved for the root node
- Thus a domain name is the sequence of labels from a node to the root, separated by dots (“.”s), read left to right
  - name space has a maximum depth of 127 levels
  - domain names are limited to 255 characters in length
- A node’s domain name identifies its position in the name space
- Traditional top level domain names are (generic three letters)
  - .mil., .gov., .edu., .net., .com., .org. each with a specific meaning (military, governmental, education, network infrastructure, (nonprofit) organizations, corporations)
- Country domains (two letters in ISO standard 3166)

Concepts – Domain Name Wars

- Explosive growth the Internet lead to growth of domain name space two
  - e.g. com and de domains are biggest toplevel domains with more the 2 million entries each
- As introduced the three letter endings had a certain meaning, but this is mostly obsoleted
  - you will find many corporations with more than one top level domains: ibm.com,net,org,us,de,... so the original idea of name space distribution is lost ...
  - most of the multi entries are redirectors
  - typical solution now to find: one main top level domain like wikipedia.org and national versions via subdomains like en,de,....wikipedia.org
- Lots of law suits filed in the beginning years of the Internet over DNS issues (name clashes, private persons vs. corporations, fraught, ...)

Concepts – Domain Name Assignments

- The resultant controversy caused the US Government (Dept. of Commerce) to take a much more active role
  - official governmental policy (the White Paper) on Internet resource administration created
- That policy resulted in the creation of ICANN
  - in the beginning: non profit organization (partly) with elected members
  - election procedure was revoked
- Main task: Decide on new top level domain labels, e.g. introduced
  - .name., .info.,.biz., ...
  - .eu., .asia., ... top levels ...

Concepts – DNS and ICANN

- Role of ICANN is to oversee administer Internet resources including
  - Addresses
    - Delegating blocks of addresses to the regional registries
  - Protocol identifiers and parameters
    - Allocating port numbers, OIDs, etc.
  - Names
    - Administration of the root zone file
    - Oversight of the operation of the root name servers
- Most important: ICANN oversees modification of the zone file that makes up the Internet DNS root

Concepts - Delegation

- Administrators can create subdomains to group hosts
- According to geography, organizational affiliation or any other criterion
- An administrator of a domain can delegate responsibility for managing a subdomain to someone else
  - But this isn’t required
- The parent domain retains links to the delegated subdomain
- The parent domain “remembers” who it delegated the subdomain to

Concept – Zones and Delegations

- Zones are “administrative spaces”
- Zone administrators are responsible for portion of a domain’s name space
  - authority is delegated from a parent and to a child

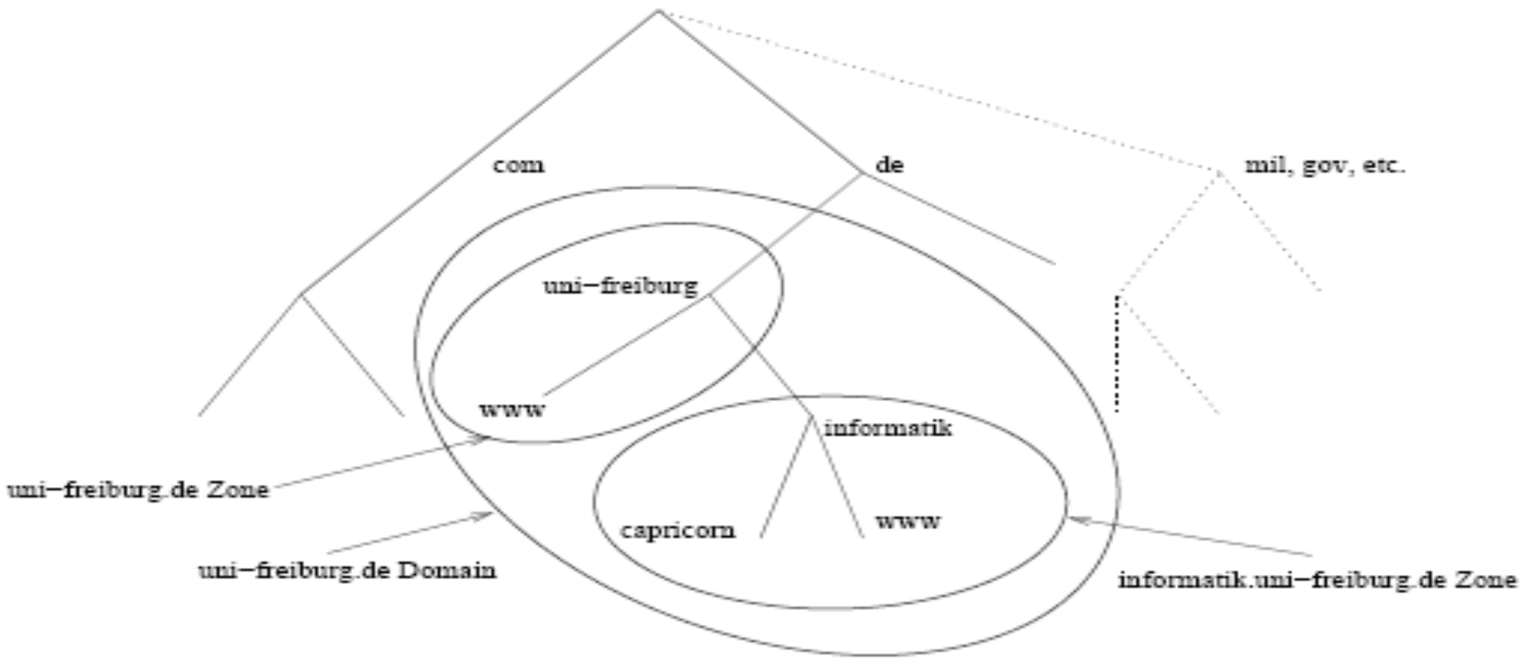

Concept – Delegations and "Forwards"

- DNS "Forward"
  - Generally, where the A records (few slides later) are
  - "Domain Names" obtained from a parent zone
  - registrar if .com, .biz, .org., and some others
  - registry if a country code (DENIC in Frankfurt for de.)
  - another organization in other cases
- Contractual - outside organization
- Formal - another part of a large organization
- Informal - from yourself to yourself

Concept – DNS Hierarchy

- The DNS imposes no constraints on how the DNS hierarchy is implemented except
  - A single root – point of vulnerability: if root nameservers are exchanged the view on data might be completely different
  - The label restrictions
- If a site is not connected to the Internet, it can use any domain hierarchy it chooses
  - Can make up whatever TLDs you want
- Connecting to the Internet implies use of the existing DNS hierarchy

Operating the database - Name Servers

- From the idea and protocol (last lecture) to the infrastructure
- Name servers answer ‘DNS’ questions.
- Several types of name servers
  - authoritative servers
  - master (primary)
  - slave (secondary)
- (Caching) recursive servers
  - also caching forwarders
  - mixture of functionality

Name Servers - Conceptual

- Authoritative
  - Give authoritative answers for one or more zones.
  - The master server normally loads the data from a zone file
  - A slave server normally replicates the data from the master via a zone transfer
- Recursive
  - Recursive servers do the actual lookups; they ask questions to the DNS on behalf of the clients
  - Answers are obtained from authoritative servers but the answers forwarded to the clients are marked as not authoritative
  - Answers are stored for future reference in the cache

Name Servers - Implementation

- Primary DNS Server (often called master)
  - maintains the master zone information
  - all changes to the information of the domain take places here
  - get propagated to the secondary servers at the Refresh interval
- Secondary DNS Server (often slave)
  - backs up the primary DNS server for a zone
  - more than one possible
- Caching
  - typically DNS of dial-in providers
    - (DSL, cable, WLAN, GPRS/UTMS, ISDN, ...)
  - improve efficiency (traffic reduction not really relevant)
  - DNS servers add answers (for a certain amount of time) from other servers to their memory

Resolver – the DNS Client

- DNS operates in classical client-server-model

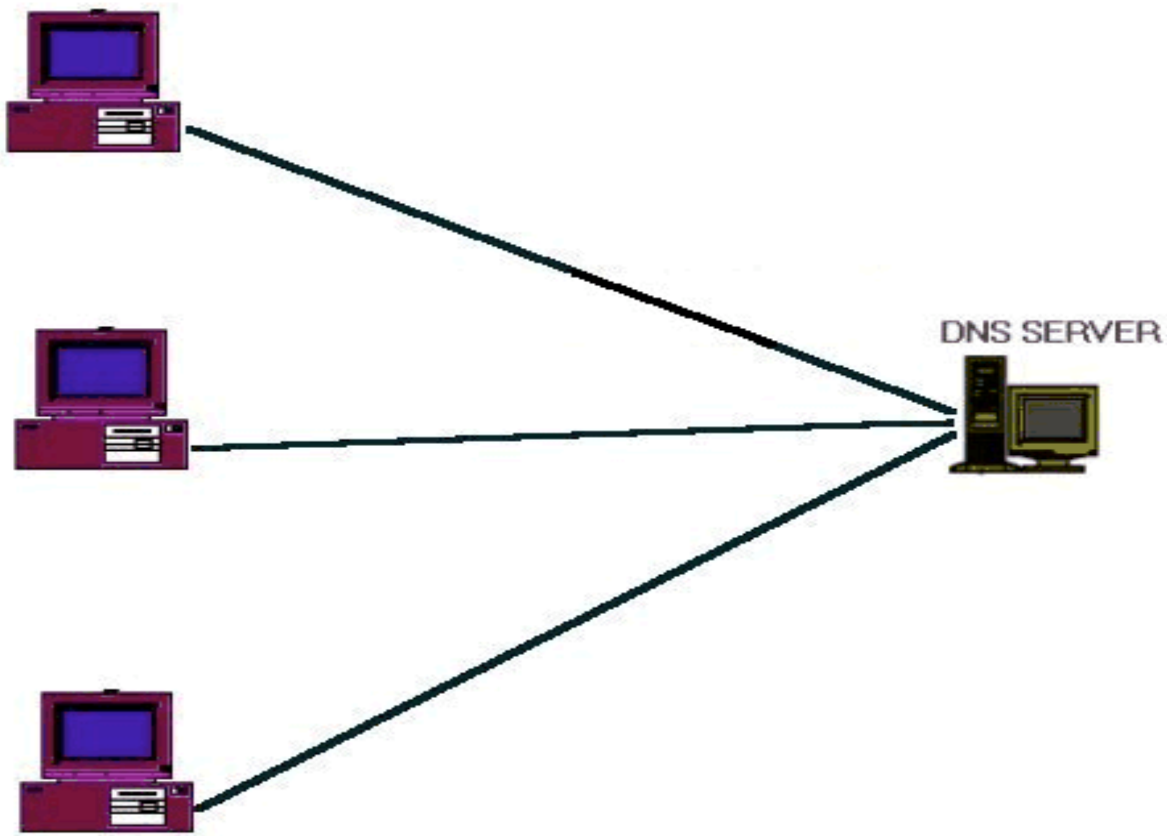

Recursion vs. Iteration

- Resolvers ask the questions to the DNS system on behalf of the application
  - asked server typically uses recursion

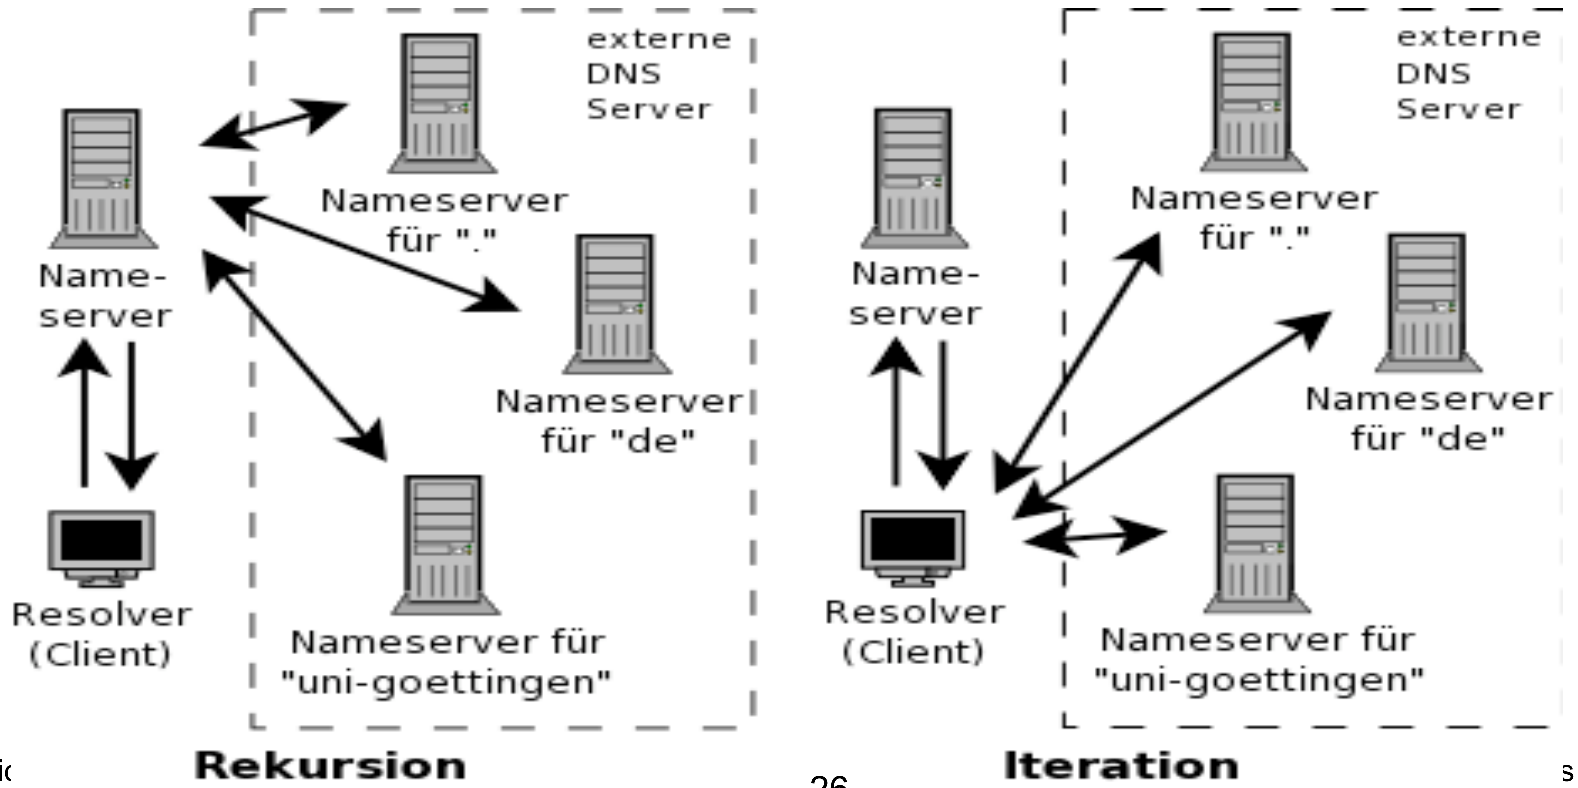

Name Resolution

- Name resolution is the process by which resolvers and name servers cooperate to find data in the name space
- To find information anywhere in the name space, a name server only needs the names and IP addresses of the name servers for the root zone (the “root name servers”)
  - The root name servers know about the top-level zones and can tell name servers whom to contact for all TLDs
- A DNS query has three parameters
  - A domain name (e.g., www.ks.uni-freiburg.de),
    - Remember, every node has a domain name!
  - A class (e.g., IN), and
  - A type (e.g., A)

Resolver – the DNS Client

- DNS clients that access name servers
  - Query name server
  - Interpret response
  - Return the information to the program requesting it
- Users do not interface directly with a DNS resolver
- Normally implemented in a system library (e.g, libc)
- gethostbyname(char *name);
- gethostbyaddr(char *addr, int len, type);

Resource Records (Basic Set)

Function	Name	Explanation
Address Records	A	Map hostname to IPv4 address e.g. *www.unifreiburg.de. IN A 132.230.6.75*
Canonical Name Records	CNAME	make one domain name an alias of another e.g. *www.uni-freiburg.de. IN CNAME www.ruf.uni-freiburg.de.*
Mail Exchange Records	MX	Specify the mail server in the domain e.g. *Foobarbaz.com IN MX 10 eric.foobarbaz.com*
Pointer Records	PTR	Map IP address to host name (reverse resolution) e.g. *75.6.230.132.in-addr.arpa. IN PTR www.ruf.uni-freiburg.de*
Name Server Records	NS	state the authoritative name servers for the domain. e.g. *foobarbaz.com. IN NS draven.foobarbaz.com.*
Start Of Authority Records	SOA	Specify that the DNS server provides authoritative information about a domain.

Resource Records

- Resource records consist of it’s name, it’s TTL, it’s class, it’s type and it’s RDATA
- TTL is a timing parameter
- IN class is widest used
- There are multiple types of RR records
- The SOA and NS records are used to provide information about the DNS itself
    - provides information about the start of authority, e.g. the top of the zone
- The NS indicates where information about a given zone can be found

Resource Records (SOA)

- Provides zone wide
  - Timing parameter
  - Master server
  - Contact address
  - Version number
- net. 3600 IN SOA A.GTLD-SERVERS.net. nstld.verisign-grs.com.
  (2006021301 ; serial
  30M ; refresh
  15M ; retry
  1W ; expiry
  1D) ; neg. answ. ttl

Resource Records (NS)

- Delegation is
  - the “glue” of the DNS system
  - is done by adding NS records:
  - sub.goe.net. NS ns1.sub.goe.net.
  - sub.ripe.net NS ns2.sub.goe.net.
- How to get to ns1 and ns2… addresses needed
  - Add glue records to so that resolvers can reach ns1 and ns2
  - ns1.sub.ripe.net. A 10.0.0.1
  - ns2.sub.ripe.net. A 10.0.0.2
- Glue is ‘non-authoritative’ data (data lives on another server, as seen in Fridays exercise)

DNS support in IPv6

- Current DNS records store 32-bits IPv4 addresses. They must be upgraded to support the 128-bits IPv6 addresses.
- A new resource record type 'AAAA' is defined, to map a domain name to an IPv6 address
- Example:
  - www.ipv6.uni-muenster.de. IN CNAME tolot.ipv6.uni-muenster.de.
  - tolot.ipv6.uni-muenster.de. IN AAAA 2001:638:500:101:2e0:81ff:fe24:37c6
  - ns.join.uni-muenster.de. IN AAAA 2001:638:500:101::53
  - ns.join.uni-muenster.de. IN A 128.176.191.10

DNS Support in IPv6

- New domains IP6.INT and IP6.ARPA are defined, to map an IP v6 address to a domain name.
- An IP v6 address is represented by a sequence of nibbles (nibble string) separated every four bits by dots with the suffix “.IP6.INT” or “.IP6.ARPA”.
- Example:
    - ; $ORIGIN 0.0.5.0.8.3.6.0.1.0.0.2.ip6.int.
    - 6.0.8.3.5.b.e.f.f.f.2.0.1.0.2.0.0.0.1.0 IN PTR atlan.ipv6.uni- muenster.de.
    - 5.f.4.7.8.d.e.f.f.f.8.1.0.e.2.0.0.0.2.0 IN PTR lemy.ipv6.uni-muenster.de.
        - or
    - ; $ORIGIN 0.0.5.0.8.3.6.0.1.0.0.2.ip6.arpa.
    - 6.0.8.3.5.b.e.f.f.f.2.0.1.0.2.0.0.0.1.0 IN PTR atlan.ipv6.uni- muenster.de.
    - 5.f.4.7.8.d.e.f.f.f.8.1.0.e.2.0.0.0.2.0 IN PTR lemy.ipv6.uni-muenster.de.

DNS Support in IPv6

- Existing queries are extended to support IP v4 and IP v6
- When both ‘A’ and ‘AAAA’ records are listed in the DNS, there are three different options:
  - return only IPv6 address
  - return only IPv4 address
  - return both IPv4 and IPv6 addresses
- The selection of which address to return, or in which order to return can affect what type of IP traffic is generated
- BIND 9.X is fully IPv6 compliant
- Problem: name space fragmentation
- Not all operating systems and not all DNS servers offer IPv6 transport lookups

Timers in DNS

- TTL is a timer used in caches
  - An indication for how long the data may be reused
  - Data that is expected to be ‘stable’ can have high TTLs
- SOA timers are used for maintaining consistency between primary and secondary servers
  - might be given in seconds (integer)
  - abbreviations possible, like on slide before
    - W – Week
    - M – Minute
    - D – Day
- Because of timing issues it might take some time before the data is actually visible at the client side

DNS Extensions - ENUM

- DNS is a rather successful concept for the distribution of vital network information (mostly by now mapping names to IPs and vice versa)
- DNS can also be used to map phone numbers to URIs
- Addressing (naming) on the Internet:
  - IP addresses: 132.230.121.6
  - domain names: www.ks.uni-freiburg.de
  - Uniform Resource Identifiers (URIs)
    - mailto: dsuchod@rz.uni-freiburg.de
    - http://132.230.6.72
    - http:/www.ks.uni-freiburg.de
    - sip:dirk@siphone.de

DNS - ENUM

- Voice-over-IP is an emerging trend for some years
  - problem: how to merge the totally different numbering schemes in the IP and telephony world
- Addressing (numbering) on the PSTN:
  - E.164 “phone” numbers: +49 761 203 4698
- Why telephone numbers any more?
  - people know how to use phone numbers
  - billions of devices only use numeric key pads, especially wireless devices
  - many VoIP customers use normal phones with terminal adapters or IP phones with numeric keypads

DNS – ENUM - Definition

- Why telephone numbers any more?
  - URIs like sip:user@domain have advantages and disadvantages
  - on of their biggest problems: they cannot be dialed on the PSTN
  - Phone numbers may be used for other services on the Internet (Instant Messaging, Video, …)
  - URI's and telephone numbers will co-exist for the indefinite future
- So Electronic or E.164 NUMber mapping is defined by the Internet Engineering Task Force (IETF) in RFC3761

DNS – ENUM – e164.arpa tree

- The e164.arpa domain was selected by the Internet Architecture Board specifically for this purpose with the concurrence of the ITU
- .ARPA is designated by the IAB for Internet Infrastructure issues
  - in-addr.arpa (reverse IP address look up)
- .ARPA is a well managed, stable and secure operational environment under IAB supervision
- Single domain structure under e164.arpa becomes the authoritative “root” for E.164 telephone numbers

DNS – e164.arpa tree - Tiers

- ETSI (European Telephone Standardization Institute) defines so called Tier level
  - Tier-0 - The registry operator for e164.arpa and its name servers
  - Tier-1 - Registry for a “country”: e.g. 4.4.e164.arpa
  - Codes are not just for countries: satellite operators, multinational telcos, international free phone numbers
  - Tier-2 - Registrars who process registration requests
  - Not area code level delegations as the terminology might suggest
- Problems would occur if alternate trees are operated ...

DNS – ENUM

- Why DNS and not some other Internet service?
- DNS
  - It’s there …
  - It works…
  - It’s global…
  - It scales…
  - It’s open…
  - Anyone can use it…

ENUM – Major Benefits

- The mapping of „Telephone Numbers“ to Uniform Resource Identifiers (URIs) using the Domain Name System (DNS) in the domain e164.arpa
  - URIs are used to identify resources on the Internet (e.g. http://enum.nic.at)
  - The purpose of ENUM is to enable the convergence between the PSTN and the Internet
- ENUM can be used for any URI = any service
  - mailto, fax, video, …
  - sms, mms, …
  - h323, pres, im, …
  - http, ftp, certificates, locations, …

ENUM – Concepts

- ENUM should not be mistaken for:
  - A real-time call forwarding service
  - ENUM should not be used to implement a follow-me service, modifying ENUM entries in real-time depending on location, time-of-day, etc.
  - This should be done as a SIP service at the SIP proxy (later lectures)
  - A „presence“ service - presence should also be implemented at the SIP proxy (e.g. with SIMPLE)
  - ENUM does not provide NOTIFY and also no policies
  - But ENUM may point to a presence service or to a geo location, e.g. for a company or a hotel

ENUM – DNS Mapping

- take an E.164 phone number
- +49 761 203 46 98
- remove the “+”, spaces and other non cipher characters
- turn it into a FQDN
- 8.9.6.4.3.0.2.1.6.7.9.4.e164.arpa.
- returns list of URIs
- sip:dirk@siphone.de
- query the DNS (for NAPTR)
- mailto:dsuchod@rz.uni-freiburg.de
- sms tel:+497612034698

DNS – New Record Type - NAPTR

- NAPTR - resulting name looked up in the DNS
- Horribly complex :-)
  - Define preferences and order to reach services
  - Can include regular-expressions and substitutions
  - Ultimately identify URIs
  - Example:
    - NAPTR 100 10 "u" "sip+E2U" \ "!^.*$! sip:jim@sip.uni-freiburg.de!"

DNS – New Record Type - NAPTR

- How to reach a SIP gateway for some phone number
- Order and Preference fields allow intelligent selections of services & protocols to be made:
  - “Send email if the SIP gateway is unable to process fax now”
  - “Don’t call my cellphone when I’m overseas”
  - “Divert to voicemail if busy”
- There are other extensions to DNS not handled in this course (key service for secure transactions, IDNS, ...)

ALBERT-LUDWIGS-
UNIVERSITÄT FREIBURG

Communication Systems

DNS

University of Freiburg
Computer Science
Computer Networks and Telematics
Prof. Christian Schindelhauer

CoNe Freiburg

IIF
INSTITUT FÜR INFORMATIK FREIBURG