The Domain Name System

David Conrad

David.Conrad@nominum.com

Nominum, Inc.

Overview

- Introduction
- History
- Name space structure
- Technical details
- Administrative details
- Political details
- Futures
- Summary

Introduction

- For the Internet to operate, certain globally unique identifiers must exist
  - Protocol numbers, port numbers, addresses, names, etc.
- Administration of these identifiers is done by the Internet Assigned Numbers Authority (IANA)
  - The IANA delegates the administration of some of these resources to other entities
  - Names are by far the most contentious

Names vs. Addresses

- In the Internet, an address provides information on how to reach a particular place
  - Usually hierarchical in nature
    - Cherry Hills Ogikubo #301, 4-6-6 Ogikubo Suginami-ku, Tokyo, Japan
    - +1-808-329-6085
    - 202.12.28.129
- Names identify an object once its location is known
  - Any hierarchy is administrative only
    - David R. Conrad
    - Tokyo
    - isc.org
- People use names, machines use addresses

The Domain Name System

- A system which permits humans to use names and machines to use addresses
- Scalable
  - Over 90 million entries in the global DNS now
- Consistent
  - You get the same answer where ever you ask
- Resilient
  - Specifically designed to avoid single points of failure
- Without the DNS, the Internet would not be usable

DNS in a Nutshell

- DNS is a distributed database
  - Data is maintained locally, but available globally
- DNS uses
  - replication to achieve robustness
  - caching to achieve adequate performance
- DNS is composed of
  - a namespace
    - the database's structure
  - name servers
    - store data from specific segments of the database Answer questions from...
  - resolvers
    - translate applications' requests for data into DNS queries
    - Interpret name server's responses

Overview

- Introduction
- History
- Name space structure
- Technical details
- Administrative details
- Political details
- Futures
- Summary

In the Beginning...

- There was the ARPANET's *HOSTS.TXT* file
  - *HOSTS.TXT* mapped every ARPANET host's name to its IP address
  - Format of an entry looked like:
    - HOST:<address>:<name,aliases>:<hardware>:<os>:<list of services>
    - e.g.,: HOST : 10.2.0.52 : USC-ISIF,ISIF : DEC-1090T : TOPS20 :TCP/TELNET,TCP/SMTP,TCP/FTP,TCP/FINGER,UDP/TFTP :
  - With this simple format, mapping from name to address ("forward mapping") and from address to name ("reverse mapping") is easy
    - On Unix systems, the HOSTS.TXT file was converted to /etc/hosts format

Life with HOSTS.TXT

- Easily implemented and understood
- Everybody (in theory) had the same version of the file
- The file was maintained by the SRI Network Information Center (the "NIC")
  - All file edits done by hand
- Network administrators sent updates via the net
  - Initially via electronic mail
  - Later via FTP
- The NIC released updated versions of the file twice a week

The Network Explodes

- Around 1980, the ARPANET consisted of hundreds of hosts
- The ARPANET changed networking protocols from NCP to TCP/IP
  - NCP required hardware (IMPs)
  - TCP/IP was implemented in software
    - And thanks to the U.S. government, the software was essentially free
- LANs became popular
  - And engineers figured out how to use ARPANET hosts as “routers” so that any host on the same LAN could use the ARPANET

The Problems with *HOSTS.TXT*

- Consistency
  - The network changed more quickly than the file was updated
- Name collisions
  - No two hosts could have the same name
    - “Good” names quickly exhausted
  - There was no good method to prevent duplicate names
    - Human intervention was required
- Traffic and load
  - The traffic generated by downloading the file became significant
    - Download time sometimes longer than update period
- The model didn't scale well

Solving the Problem

- ARPANET powers-that-were launched an investigation into replacement for HOSTS.TXT
- Goals:
  - To solve the problems inherent in a monolithic host table system
  - Have a consistent naming structure
  - Create a generic solution that can be used for multiple purposes

The New Naming System

- Requirements:
  - Decentralized administration
    - With data updated locally, but available globally
  - A hierarchical name space
    - To guarantee unique names
  - Massive scalability
- Assumptions:
  - Database size will be proportional to the number of users, not hosts
  - Names are long term persistent

The Advent of DNS

- Paul Mockapetris, then of USC's Information Sciences Institute, designed the architecture of the new system, called the *Domain Name System,* or *DNS*
- The initial DNS RFCs were released in 1984:
  - RFC 882, “Domain Names - Concepts and Facilities”
  - RFC 883, “Domain Names - Implementation and Specification”
- The transition plan was initially released in November, 1983, transition to be completed by May, 1984

The DNS RFCs

- RFCs 882 and 883 were superseded by:
  - RFC 1032, “Domain Administrators Guide”
  - RFC 1033, “Domain Administrators Operations Guide”
  - RFC 1034, “Domain Names -- Concepts and Facilities”
  - RFC 1035, “Domain Names -- Implementation and Specification”
- Additional RFCs specified
  - New “resource record” types
  - DNS operational considerations
  - DNS policies
- DNS continues to evolve to meet the changing demands of the Internet

Overview

- Introduction
- History
- Name space structure
- Technical details
- Administrative details
- Political details
- Futures
- Summary

The Name Space

- The *name space* is the structure of the DNS database
- It's an inverted tree of nodes with the **root** at the top
- Each node has a **label**
- The root node has a null label, written as "."

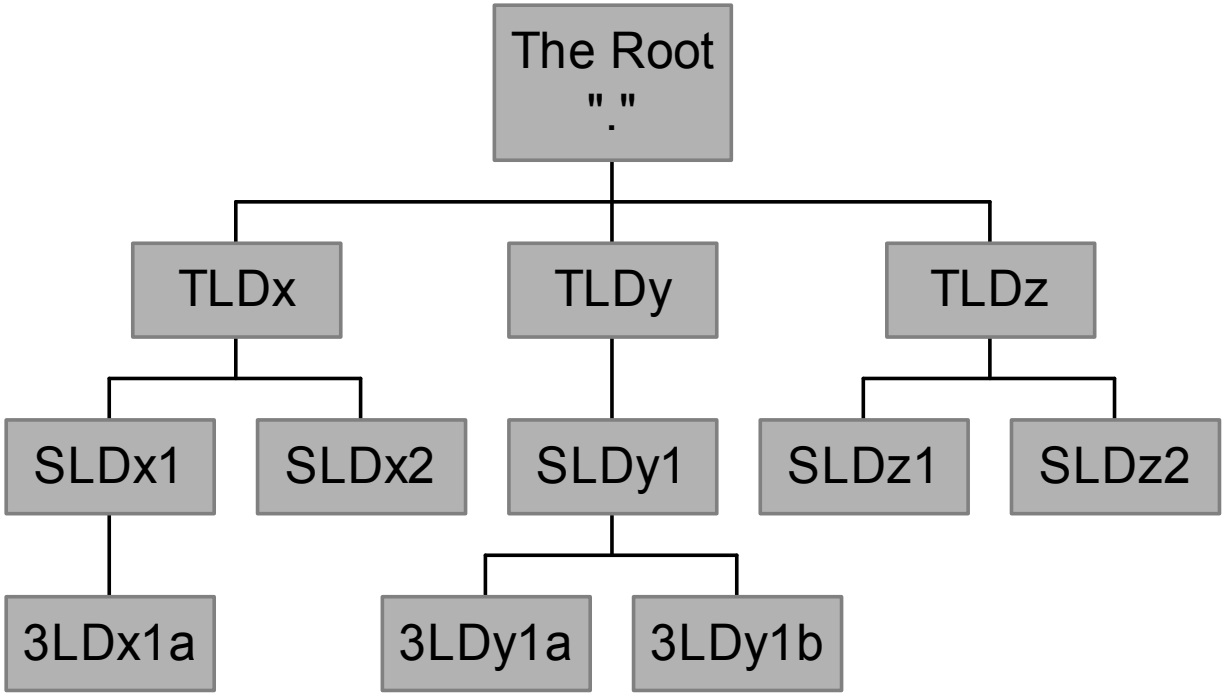

Domains

- A *domain* is a node in the name space and all its descendants
  - That is, a subtree of the name space
- A domain's domain name is the same as the name of the node at the root (top) of the subtree

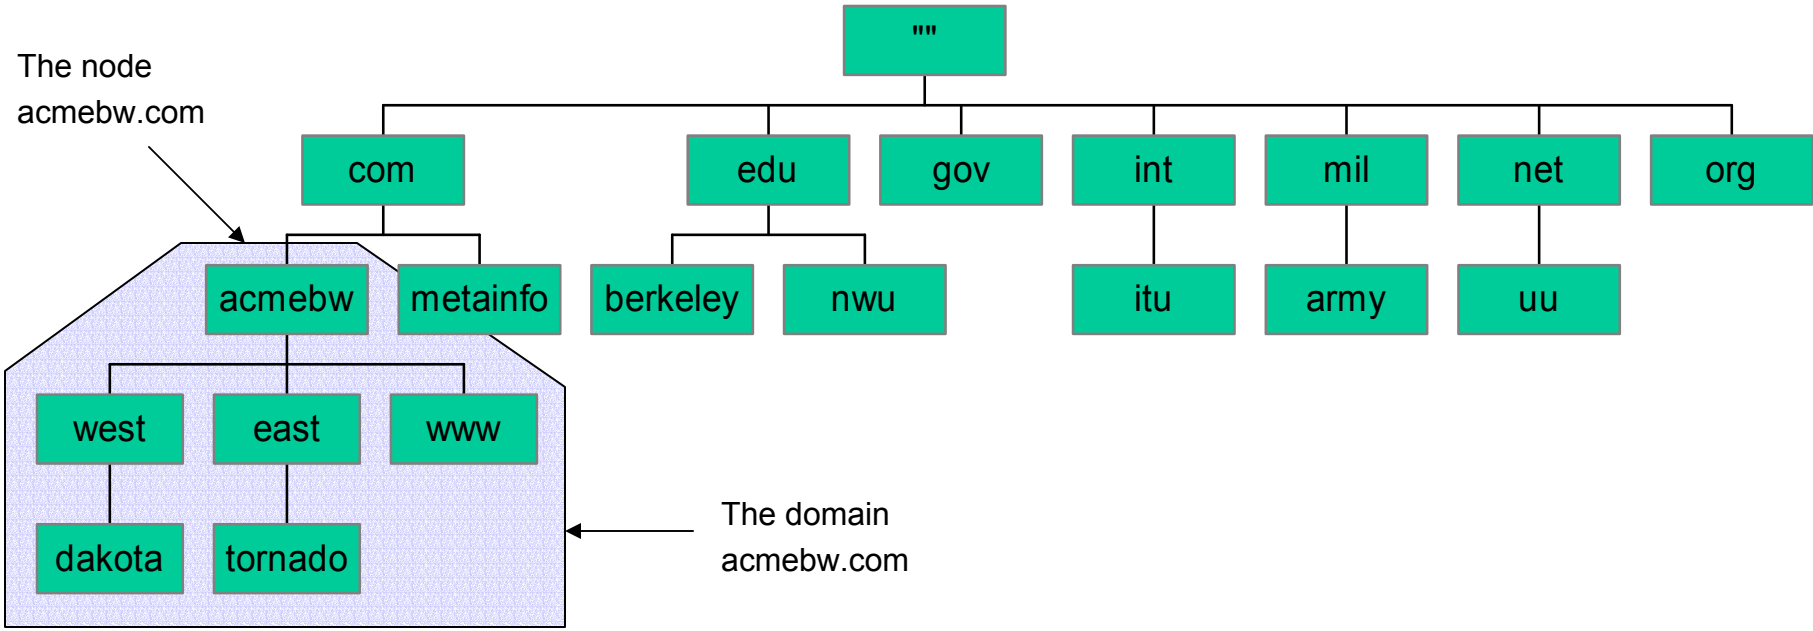

Subdomains

- One domain is a *subdomain* of another if its root node is a descendant of the other's root node
- More simply, one domain is a subdomain of another if its domain name ends in the other's domain name
  - So *sales.acmebw.com* is a subdomain of *acmebw.com*
    - Also of .*com*, but that isn't usually stated
  - *acmebw.com* is a subdomain of *com*

The Levels

- Above the top is the root.
- Beneath the root are the “Top Level Domains”
  - e.g., .COM, .JP, .INT, etc.
- Beneath the Top Level Domains are “Second Level Domains”
  - e.g., Nominum.COM, AD.JP, ITU.INT
- Beneath the Second Level Domains are “Third Level Domains”
  - e.g., www.Nominum.COM, IIJ.AD.JP, www.ITU.INT
- And so on...

The Root

- The DNS provides a coherent, consistent namespace via a **singly** rooted hierarchical tree structure
  - This root holds the definition of all top level domains that are guaranteed to be unique in that DNS tree
- THERE CAN ONLY BE ONE!
  - Violation of this rule results in inconsistencies in the namespace
    - That is, a name can translate to different addresses depending on where you ask the question
- Due to protocol limitations there are 13 nameservers that serve the root zone
  - a-m.root-server.net
    - a.root-server.net is the primary
- The root nameservers are provided in a configuration file
  - Control of this file is becoming an issue

TLD Structure

- In 1983 (RFC 881), the idea was to have TLDs correspond to network service providers
  - e.g., .ARPA, .DDN, .CSNET, etc.
    - Bad idea -- if your network changes, your email address changes
- By October, 1984 (RFC 920), the concept of functional domains (e.g., .GOV for Government, .COM for commercial, .EDU for education, etc.) was established
  - "The motivation is to provide an organization name that is free of undesirable semantics."
- RFC 920 also provided for
  - Country domains
  - "Multiorganizations"
    - large, composed of other (particularly international) organizations
- The RFC 920 TLD structure remained stable until 1997 or so
  - More on this later...

The .ARPA Hack

- The DNS provides obvious and elegant name to address mapping
- The reverse (address to name) is a bit less elegant
  - Create a domain out of the dotted quad IP address
  - Reverse the ordering to allow for proper delegation
  - Create a “special” domain to hold the delegations
- For Example:
  - 5.10.8.128.in-addr.arpa → umd5.umd.edu
- Originally, a IN-ADDR top level domain was used
  - This was felt inappropriate, so in-addr was moved under .ARPA
- This technique has some problems when dealing with non 8-bit aligned IP address blocks

Country Domains

- With RFC 920, the concept of domains delegated on the basis of nations was recognized
- Conveniently, ISO has a list of “official” country code abbreviations
- The IANA likes using lists others define
  - Can always blame someone else…
- The ISO 3166 list is officially available from:
  - `http://www.din.de/gremien/nas/nabd/iso3166ma/codlstp1.html`
- The IANA also uses International Postal Codes for country domains
- In either case, the IANA has no control over what is in those lists

Country Domains (cont'd)

- How each country top-level domain is organized is up to the country
  - Some, like Australia's *au,* follow the functional definitions
    - *com.au, edu.au,* etc.
  - Others, like Great Britain's *uk* and Japan's *jp,* divide the domain functionally but use their own abbreviations
    - *ac.uk, co.uk, ne.jp, ad.jp,* etc.
  - A few, like the United State's *us,* are largely geographical
    - *co.us, md.us,* etc.
  - Canada uses organizational scope
    - *bnr.ca* has national scope, *risq.qc.ca* has Quebec scope
  - Some are flat, that is, no hierarchy
    - *nlnet.nl*, *univ-st-etienne.fr*
- Considered a question of national sovereignty

Current TLDs

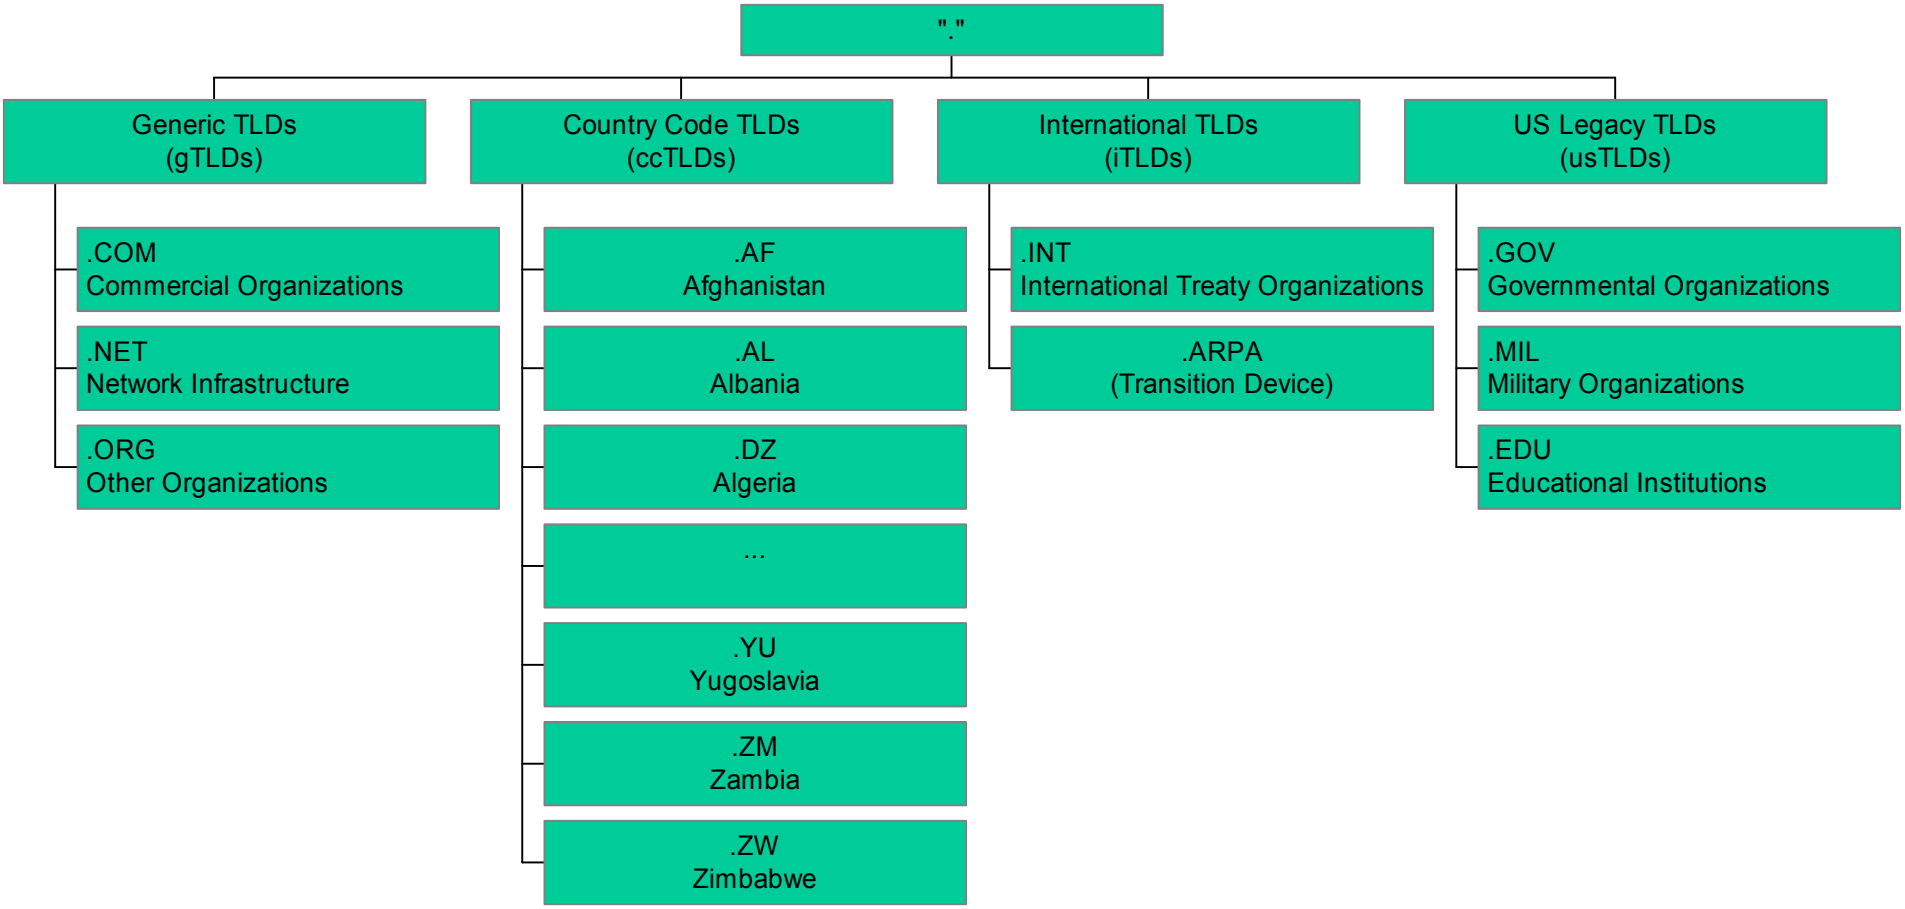

Restrictions on Labels

- The null label is reserved for the root node
  - The terminal “.” can be left off
- Labels cannot exceed 63 characters
  - Legal characters on the Internet are alpha-numeric and dash
- Sibling nodes must have unique labels

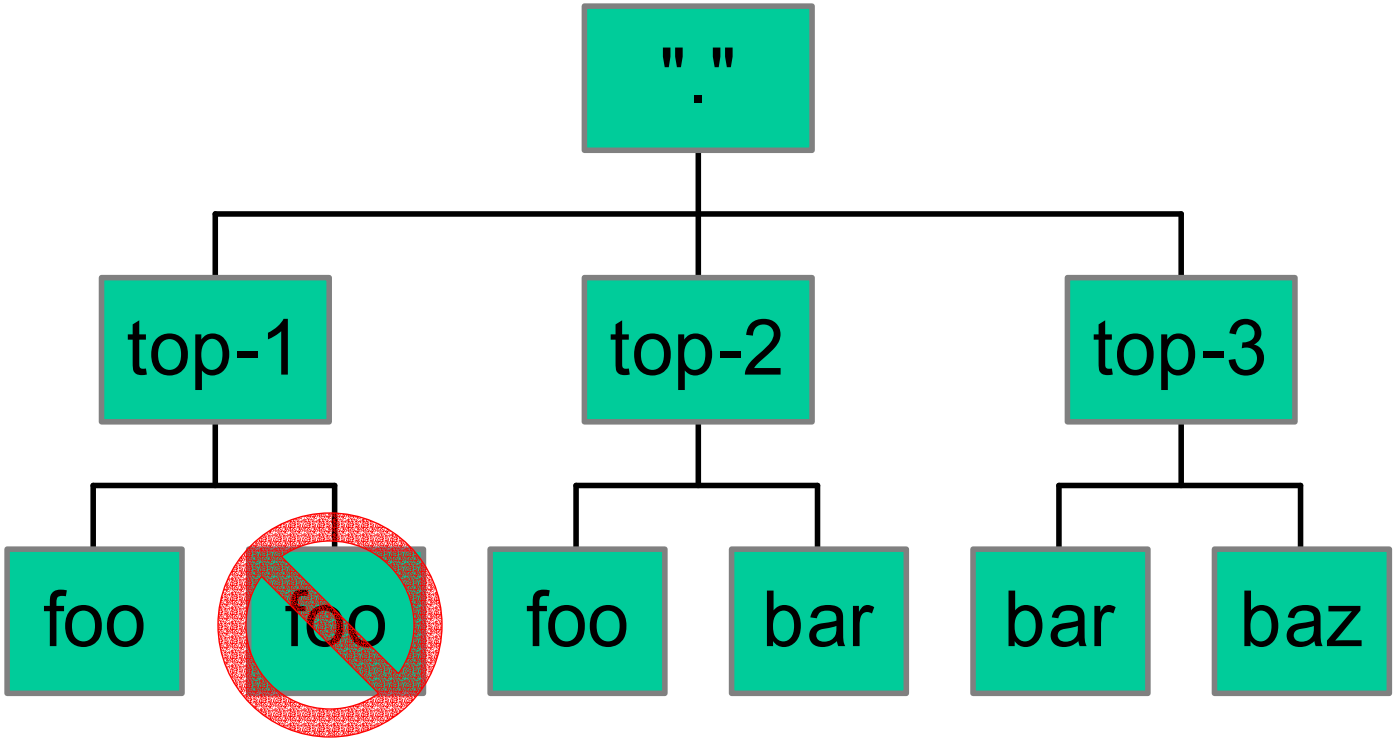

An Analogy

- The structure of the name space is similar to the many computer file systems, e.g., Unix:

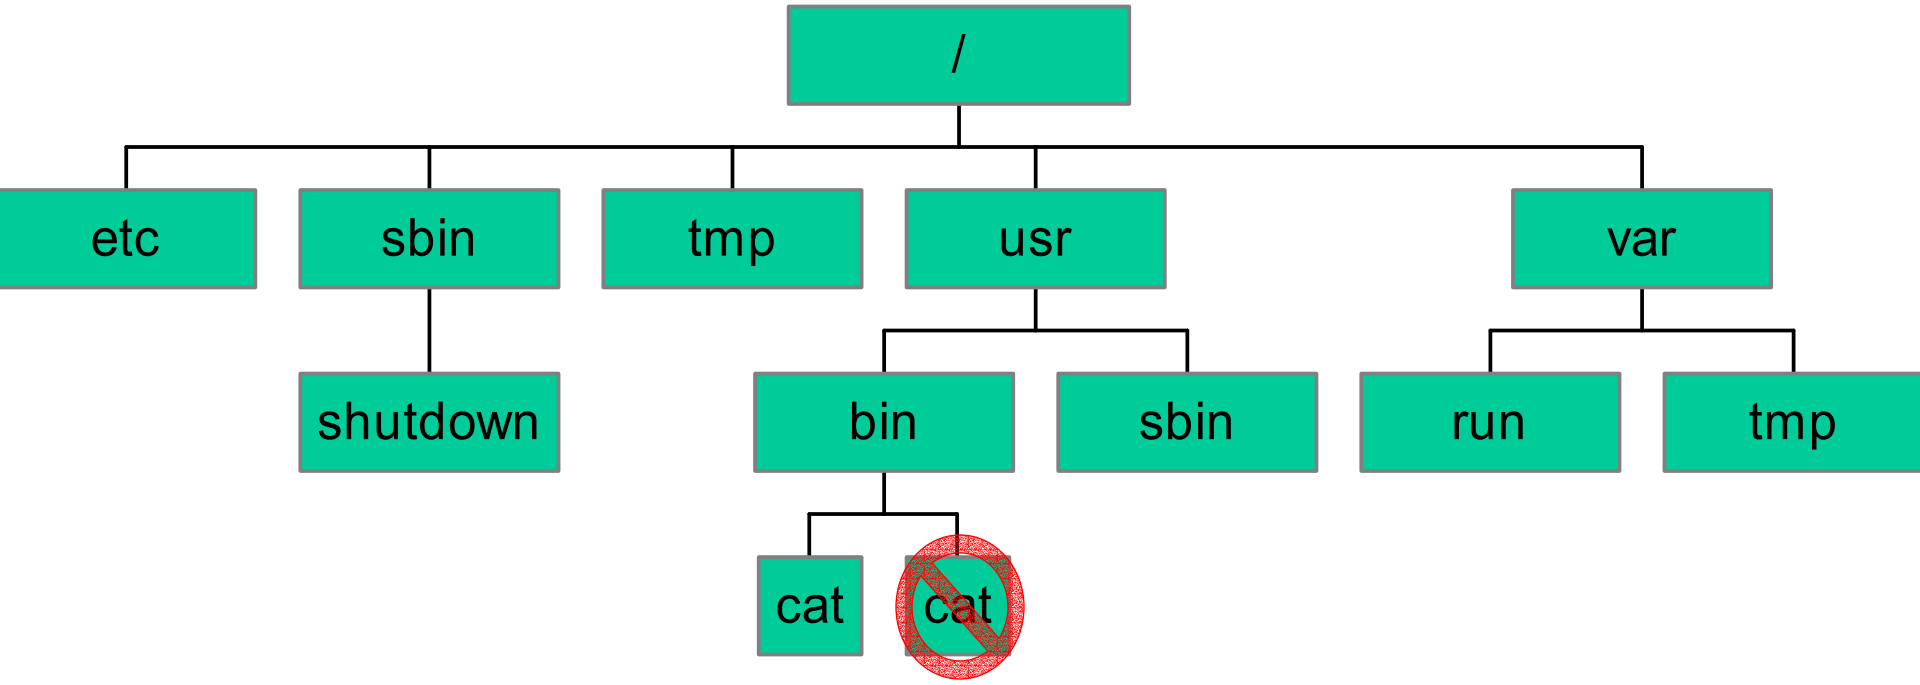

Domain Names

- A *domain name* is the sequence of labels from a node to the root, separated by “.”s
  - Each label limited to 63 characters
  - Each name limited to 255 characters
  - Maximum of 127 labels per name
- A node’s domain name identifies its position in the name space
  - Read from right (least specific) to left (most specific)
    - Similar to postal addresses in the US
      - <building> <street> <city> <state>
  - Much as a pathname uniquely identifies a file or directory in a filesystem

But What About <insert your script here>?

- RFC 952 (circa 1985) defines the Internet Host Table format (HOSTS.TXT)
  - The characters allowed in **host names** were defined as

    ```
    <name> ::= <let>[*[<let-or-digit-or-hyphen>]<let-or-digit>]
    ```
- RFC 1123 (circa 1989) relaxed legal host names to start with a number or a letter
  - RFC 1123 is the “Host Requirements” RFC
    - A standard
- Many (legacy) applications assume only RFC 1123 hostnames exist

Multi-lingual Domain Names

- RFC 1123 restrictions were enforced in BIND's resolver around 1996
  - Bad guys were putting shell meta-characters into the reverse mapping names, e.g.:
    - 129.28.12.202.in-addr.arpa →`rm -fr /*`.com
- Around 1998, people started asking why they can't have their own script in the DNS
  - Microsoft releases Win98 which permits internationalized characters in domain names
- IETF deeply concerned
  - Legacy systems have trouble with "unusual" characters

Multi-lingual Domain Names

- Today, there are at least 13 companies providing multi-lingual domain name systems
  - Most do not interoperate with each other
- The IETF has chartered the Internationalized Domain Name (IDN) working group
  - Slated to produce a specification for allowing more than [A-Za-z0-9\-] in domain names
    - Most likely by using new DNS features to transmit UTF-8 if the server can understand UTF-8, falling back to a “hostname character set” encoding if the server can’t

The Domain Name *tornado.east.acmebw.com.*

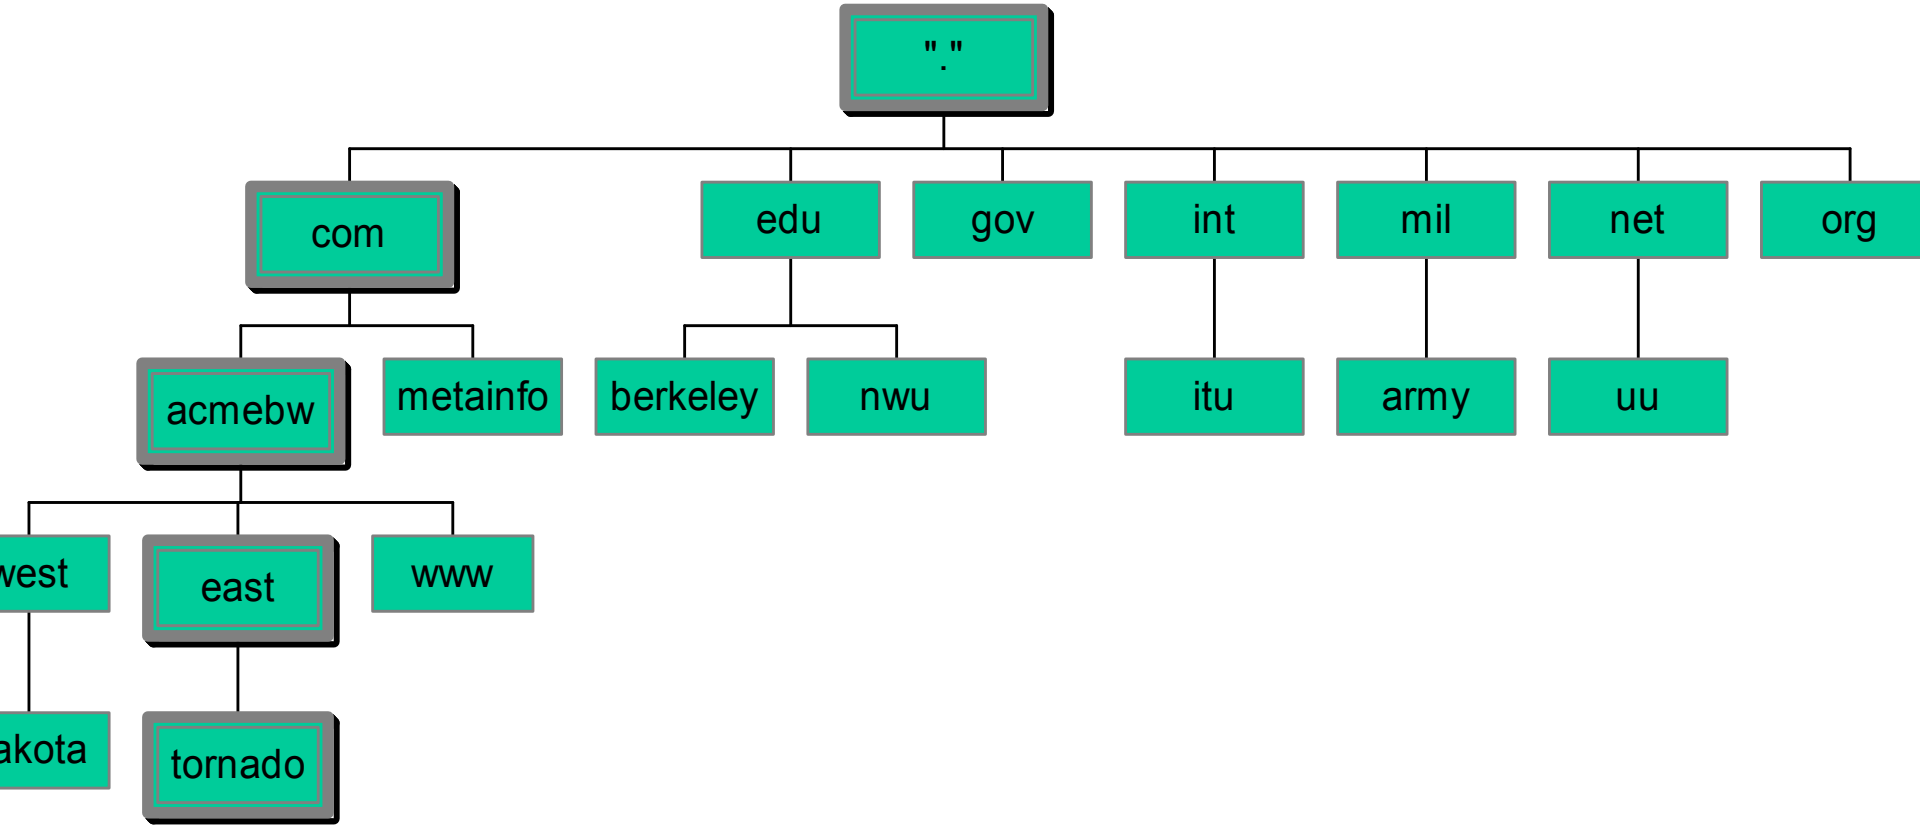

A UNIX Analogy: The File /usr/bin/cat

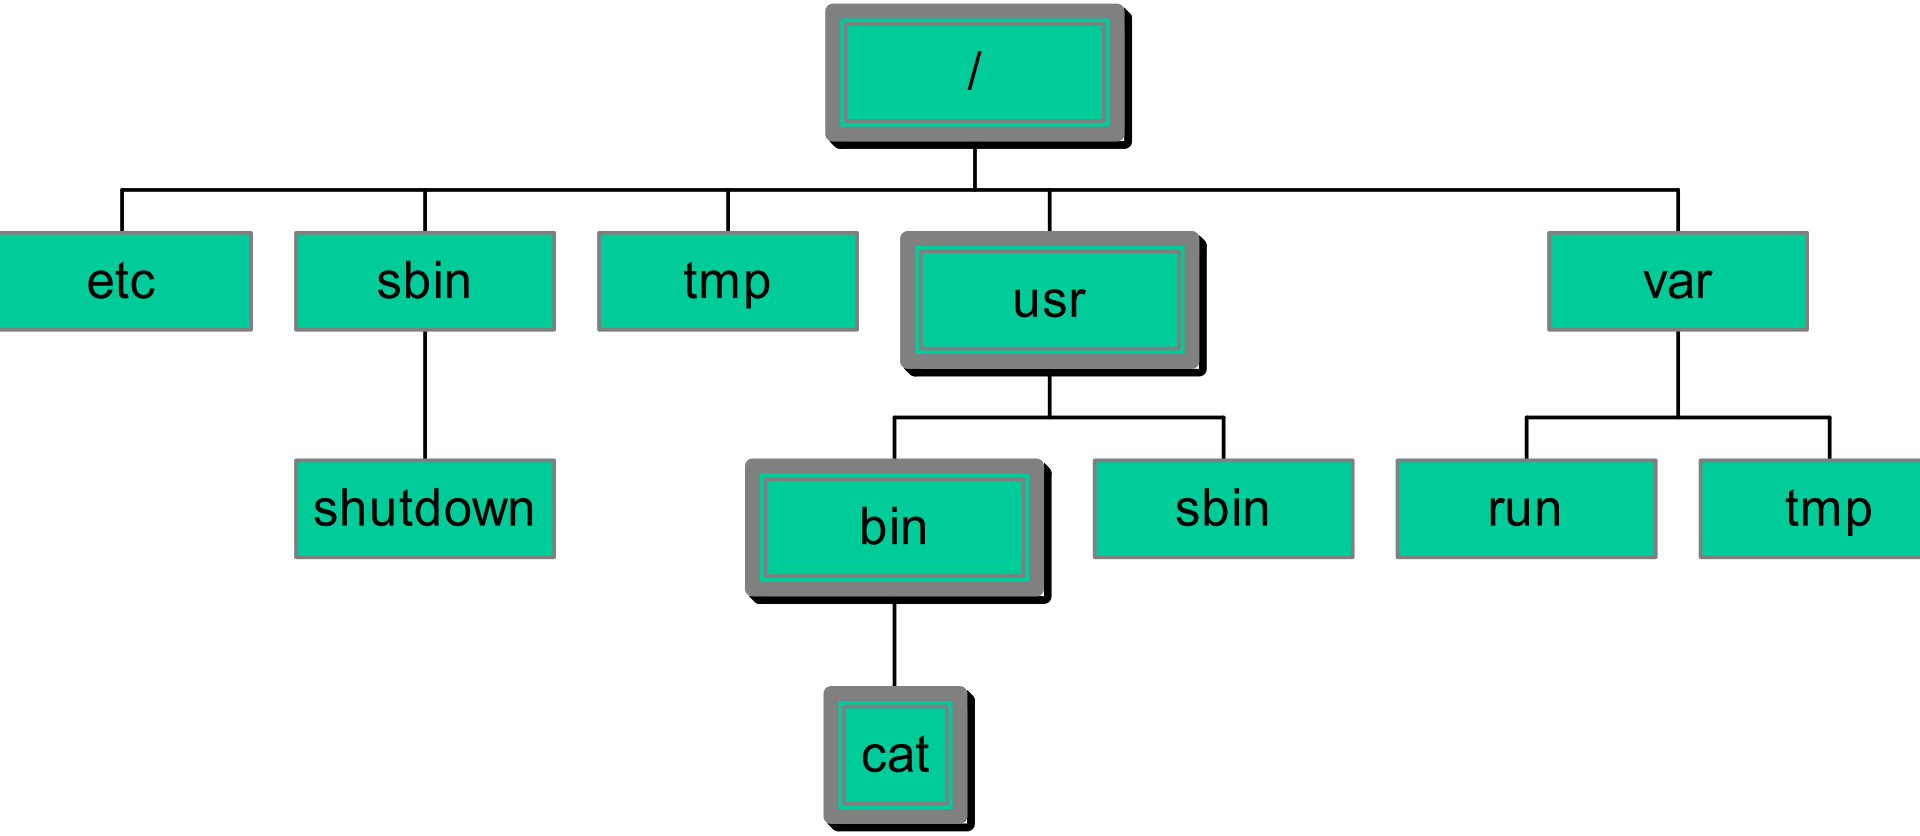

Fully-Qualified Domain Names

- A *fully-qualified domain name* (abbreviated “FQDN”) ends in a top-level domain name or a dot
  - A trailing dot (“.”) is actually the final separator between the top-level domain and the root’s null label
  - This is like absolute pathnames, which start with "/”
- Domain names without a trailing "." are not necessarily interpreted relative to the root domain
  - Just as pathnames without a leading "/" are usually interpreted relative to the current directory
- In many cases, non-absolute domain names have a domain “path” appended to them
  - Can be a security risk

Where Did the Hosts Go?

- Everywhere!
- The nodes in the name space act as indices into the distributed database
  - Some nodes represent hosts
    - These are indexes to addresses
  - Some nodes represent mail destinations
    - These are indexes to mail routing information
  - Some nodes represent an entire domain
    - These are indexes to lists of name servers
  - Some nodes are aliases for other nodes
  - A single node can represent a combination of hosts, mail destinations and domains

hp.com Domain, Host and Mail Destination

- *hp.com* is a domain (there are nodes below it)
- *hp.com* is a host in Palo Alto, California
- *hp.com* is a mail destination

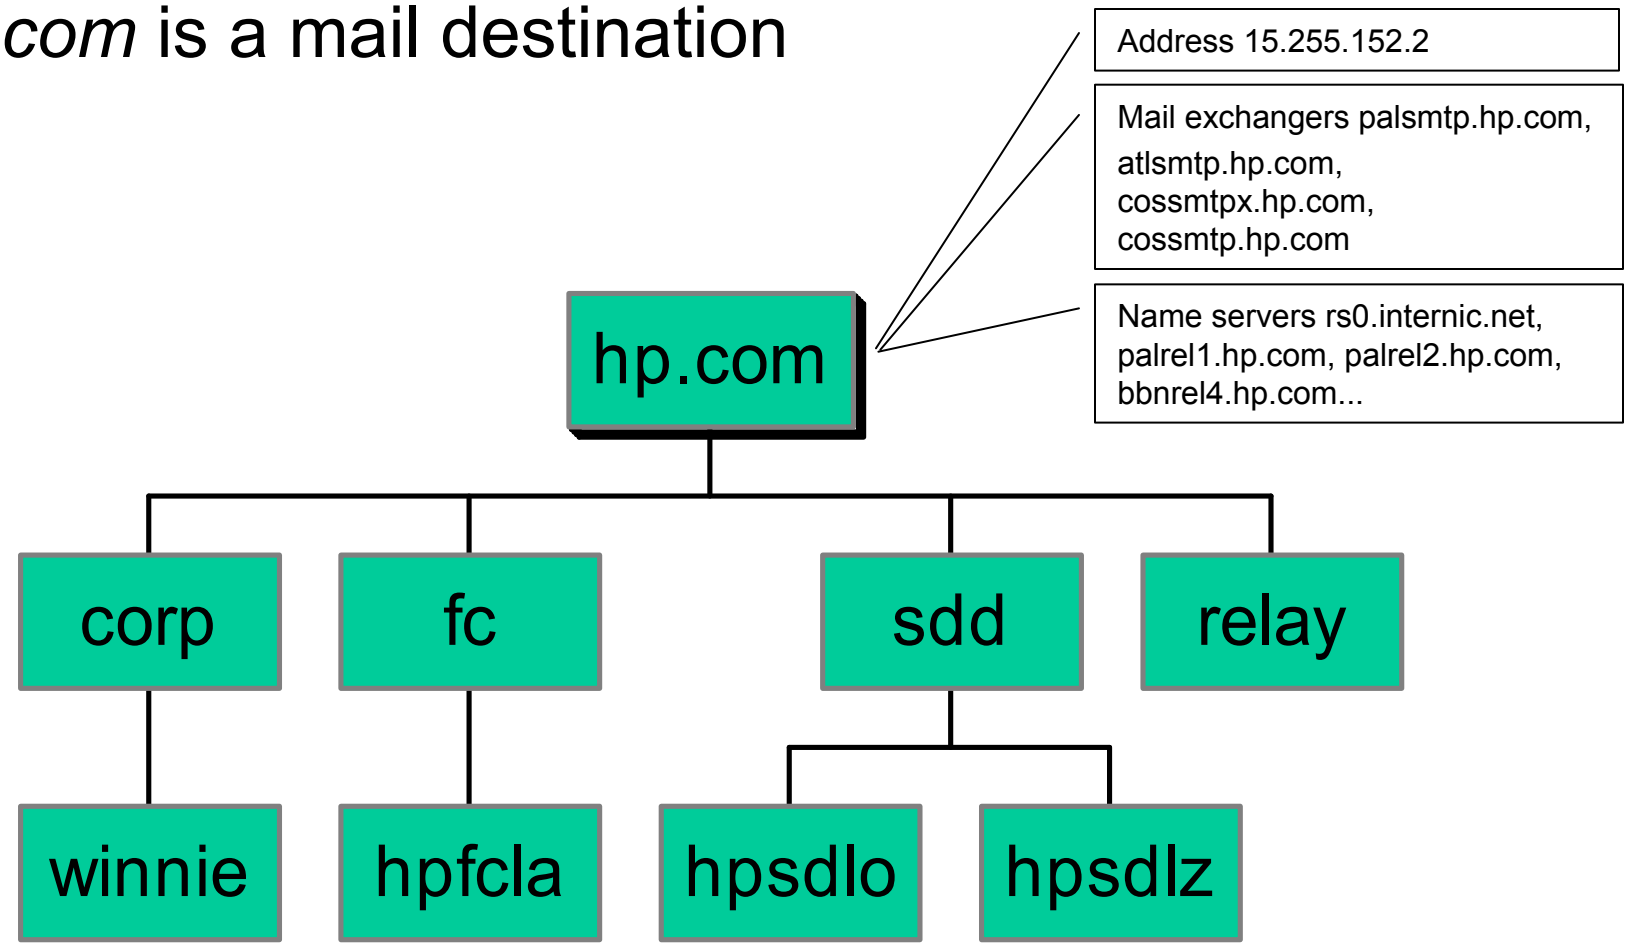

Delegation

- Administrators often create subdomains to distribute management of the domain
  - An administrator can delegate responsibility for managing a subdomain to someone else
  - The parent domain retains pointers to the sources of data for the delegated subdomain
- This sub-delegation provides for administrative scaling
  - Delegation is a good thing

Delegation Creates Zones

- Each time an administrator delegates a subdomain to someone else, this creates a new unit of administration
  - The subdomain and its parent domain can be administered independently
  - These units are called *zones*
  - The boundary between zones is a point of delegation in the name space

What's in a Zone?

- Like a domain, a zone is named after its root node
- Unlike a domain, a zone contains only descendants of the zone's root nodes that aren't in a delegated subdomain
  - Nodes below the delegation point are in another zone

into Zones

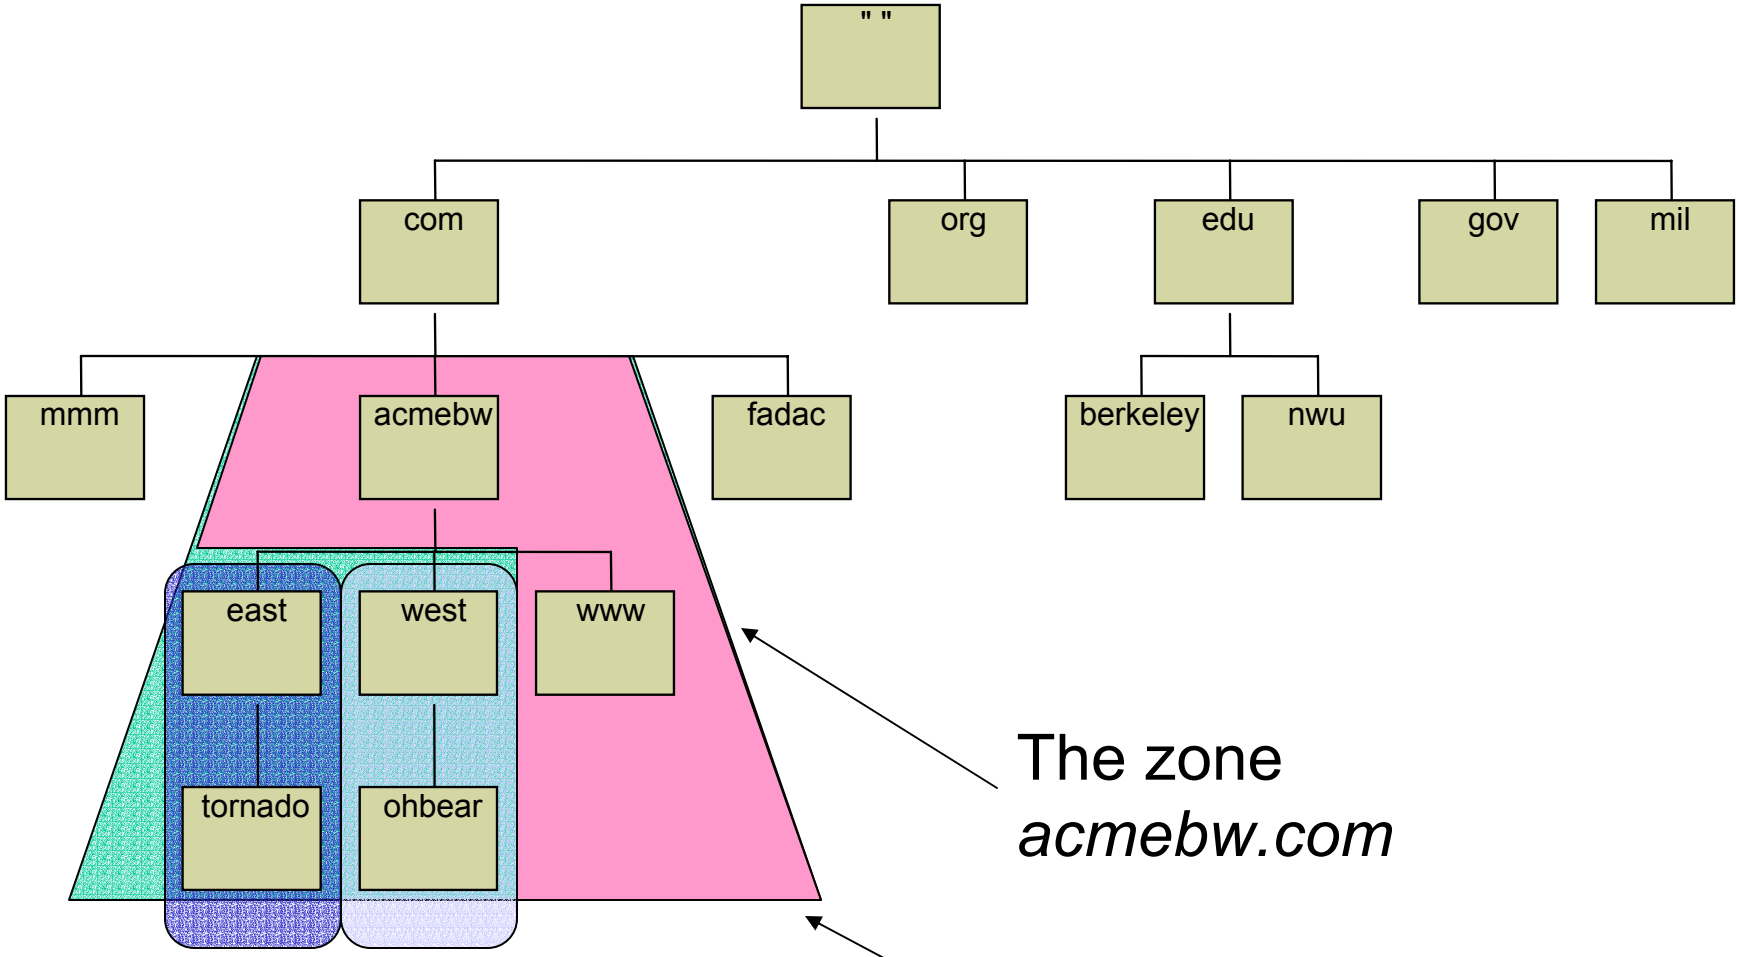
" "
com
org
edu
gov
mil
mmm
acmebw
fadac
berkeley
nwu
east
west
www
tornado
ohbear
The zone
acmebw.com
The domain
acmebw.com
The zone
east.acmebw.com
The zone
west.acmebw.com

A Delegation Example

- Think of delegation in a managerial setting
  - A manager, Rick, has overall responsibility for managing his company's internal TCP/IP network
  - However, he can't do everything himself; he delegates responsibility for some tasks to his employees
    - He delegates routing to Andy
    - He delegates email to Jeannie
    - He delegates DNS to Mike
    - But he keeps billing for himself

The "Internal" TCP/IP Network Domain Divided into Zones

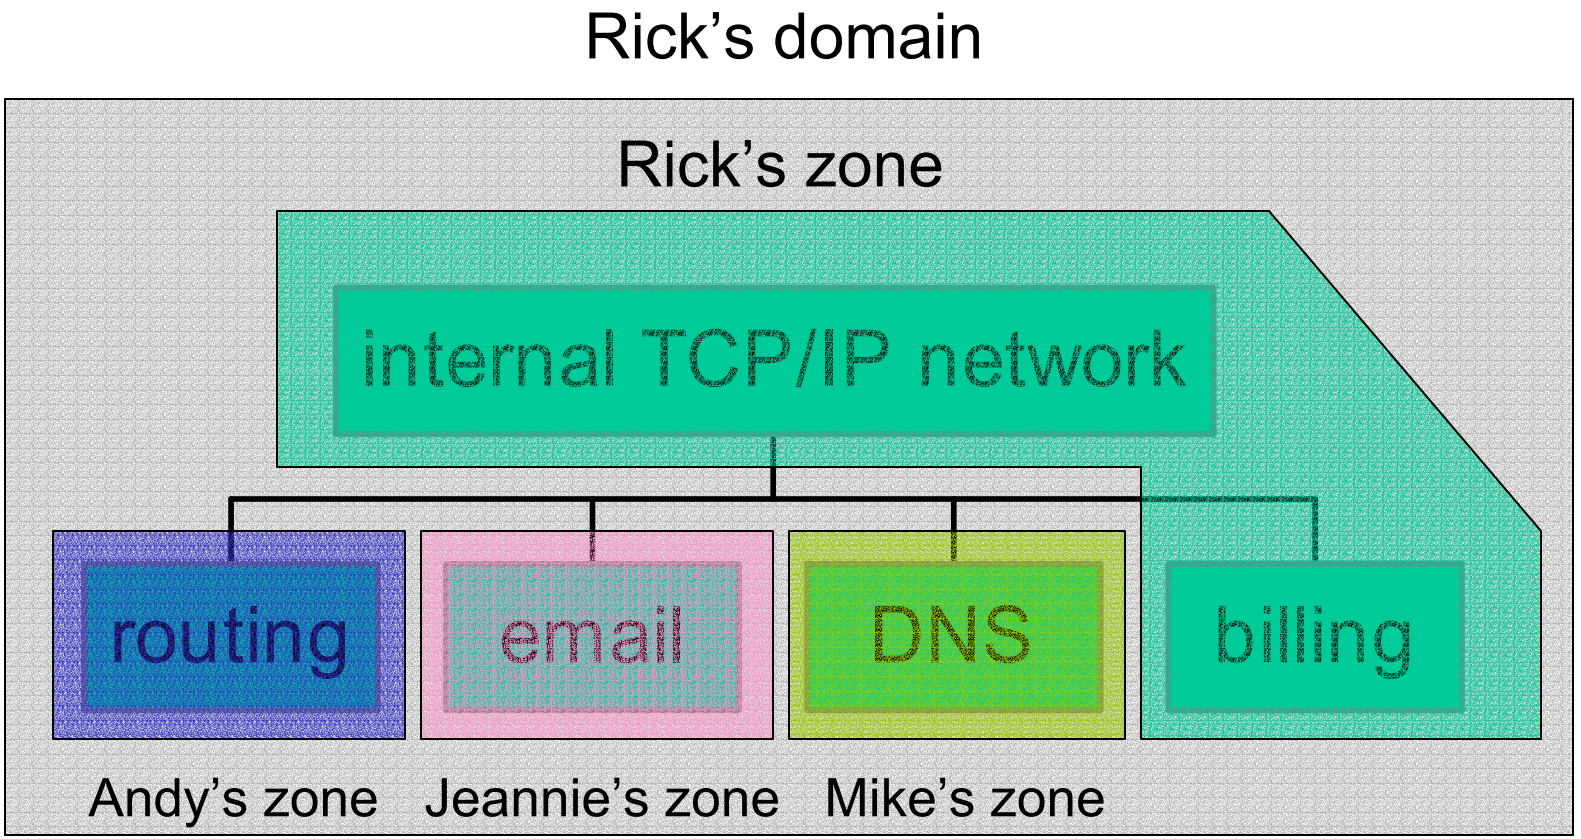

Overview

- Introduction
- History
- Name space structure
- Technical details
- Administrative details
- Political details
- Futures
- Summary

The DNS Architecture

- The Domain Name System has a client-server architecture
  - *Resolvers* are the client half
    - Always linked into an application program
      - Users execute the program, resolution requests are created and sent to servers, e.g.:
        - » netscape http://www.isc.org
      - will result in the resolver requesting the IP address(es) of www.isc.org
    - Some configuration of the resolver possible
      - Nameservers to query, timeouts, number of retries
      - On Unix, found in /etc/resolv.conf
  - *Name servers* are the server half
    - Long running server process (always active)
    - Best run on a dedicated machine
    - Resource requirements depends on many factors

Resolvers

- Resolvers are responsible for
  - *Translating* an application's request for information about a domain name into a DNS query,
  - *Sending* the query to a name server,
  - *Retransmitting* the query, if necessary,
  - *Falling back* to another name server or name service, if necessary,
  - *Translating* a name server's DNS response into a reply to the application
  - *Notifying* the application of name lookup failure (time out, authoritative non-existence, server failure, etc.)
- The operation of the resolver is almost always transparent
  - Usually a library call within an application

Name Resolution Example

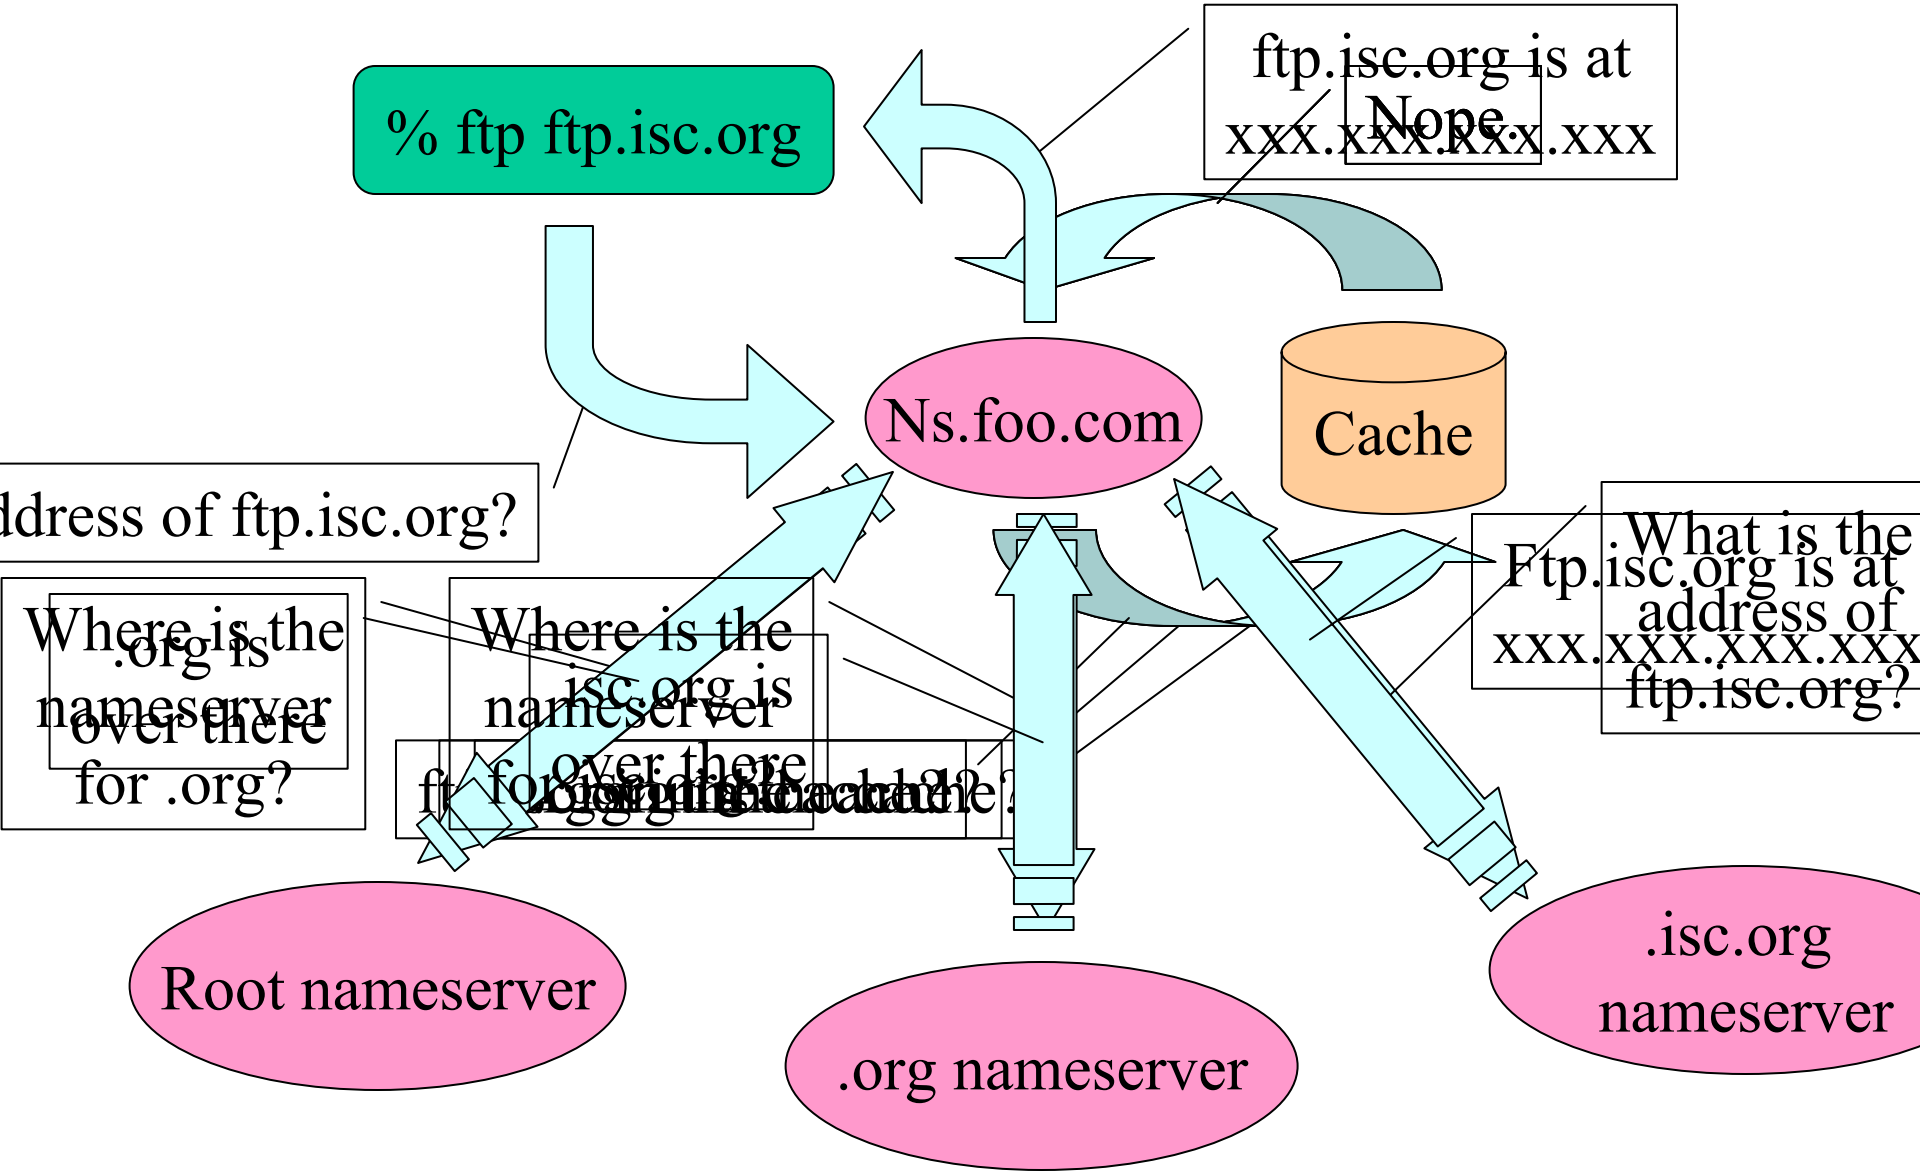

Name Servers

- Name servers are responsible for
  - *Storing* information about the name space,
    - Including additions, deletions, and changes
  - *Answering* queries from resolvers and other name servers,
  - *Querying* other name servers for information about the name space they don't already know, and
  - *Caching* information they learn about the name space from other name servers

Name Server Architecture

- You can think of a name server as part
  - *database server,* answering queries about the parts of the name space it knows about,
  - *agent,* helping resolvers and other name servers find data that other name servers know about, and
  - *cache,* temporarily storing data it learns from other name servers.

Name Server Data

- Name servers store information about the name space in units of zones
  - The name servers that load a complete zone are said to “have authority for” or “be authoritative for” the zone
- Usually, more than one name server is authoritative for the same zone
  - This ensures redundancy and spreads the loads
- Also, a single name server may be authoritative for many zones

Data in the Name Space

- Each domain name in the name space points to one or more *resource records,* or *RRs* for short
- Each resource record has a class and a type associated with it
  - The class specifies what kind of network (e.g., TCP/IP) the record describes
  - The type specifies what type of data (e.g., address) the record stores

Resource Records

- Resource records have as many as five fields, some of which are optional:
  - *Owner:* the domain name of the node to which the record is attached
  - *Time to live (TTL):* how long to keep the record in a cache
  - *Class:* the kind of network this record describes
  - *Type:* an indication of the function of this record
  - *RDATA:* record-specific data
    - The RDATA can be further subdivided into type-specific fields

Record Classes

- By far the most common class of data is the Internet class, abbreviated *IN*
  - This specifies, for example, that the addresses stored are IP addresses
- Other classes include
  - *Hesiod,* for MIT's Hesiod network protocols,
  - *CHAOSNET,* for the (largely historical) CHAOSNET protocols (also out of MIT)

Zone Data Files

- Resource records are collected into “RR Sets”
  - The collection of all RRs with the same owner
- All RR Sets associated with a zone are kept in “*zone database files*”, also called database files or db files
- Zone database files are usually named after the zone whose records they contain
  - On UNIX, usually db.zone, e.g., db.acme or db.acme.com
  - On NT with the Microsoft DNS Server, usually zone.dns, e.g., acmebw.com.dns

Name Servers and Zones

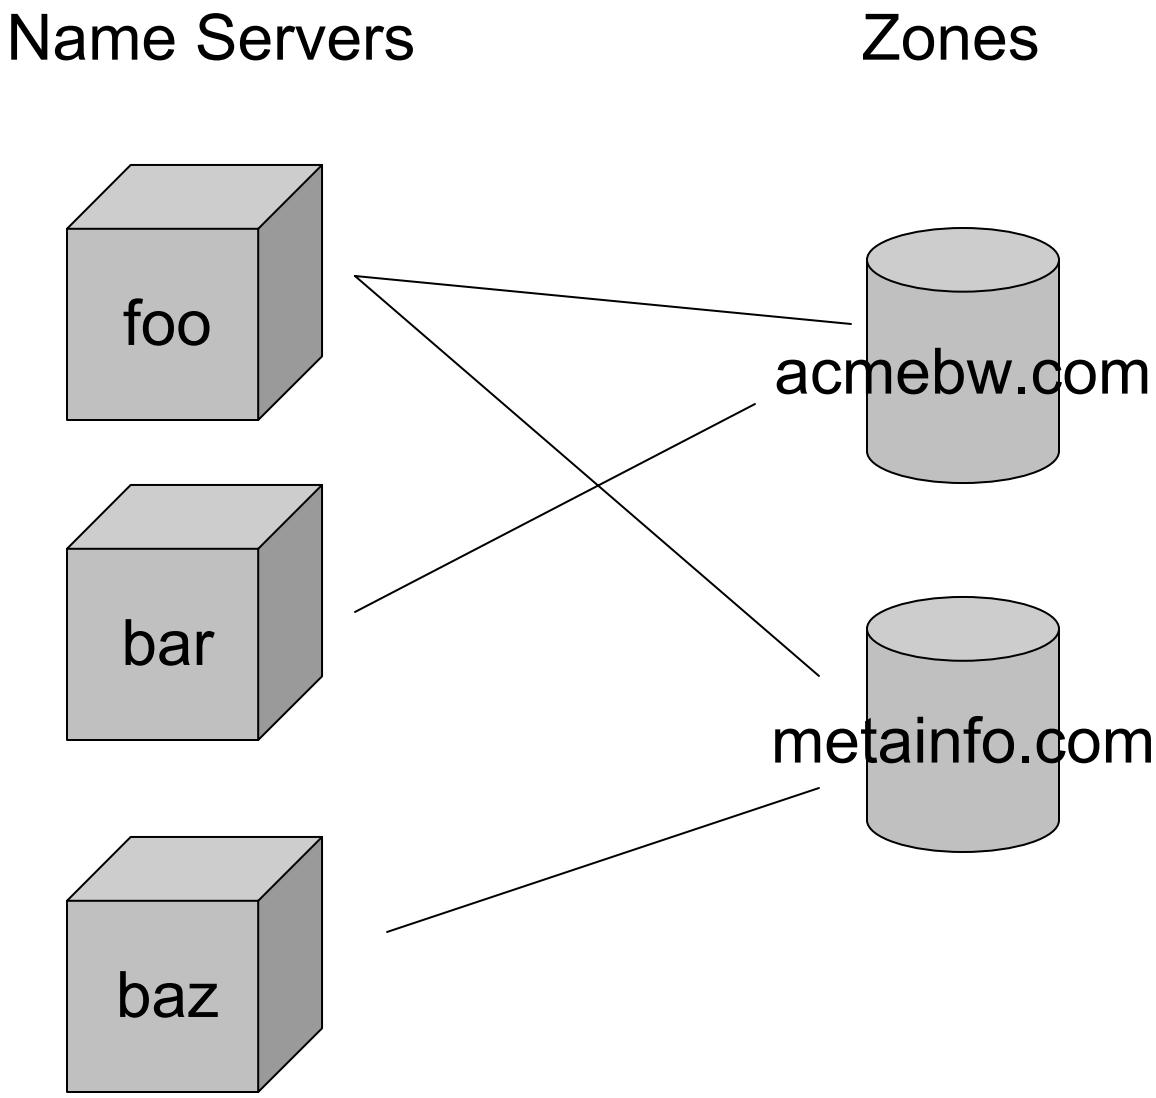

Types of Name Servers

- The *primary master* name server for a zone loads the zone's data from a file on disk
- A *slave* name server for a zone loads the zone's data from another authoritative name server (often the primary master)
  - The equivalent BIND 4 term for slave was "secondary master"
  - The server the slave gets its zone data from is called its *master* server
- A single name server can be the primary master for some zones and a slave for other zones
  - The relationship is defined zone-by-zone

Zone Transfers

- Slave servers retrieve zone data from other authoritative name servers using a *zone transfer*
- The zone transfer is initiated by the slave
  - By initiating a TCP connection to the master name server

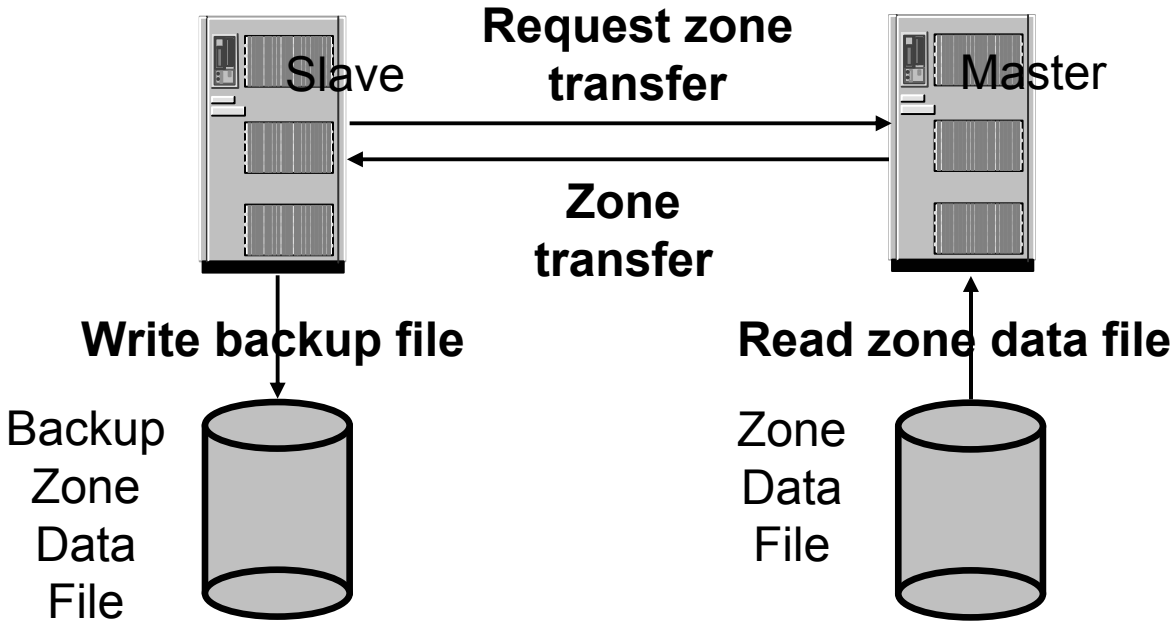

A Simple Primary Master/Slave Configuration

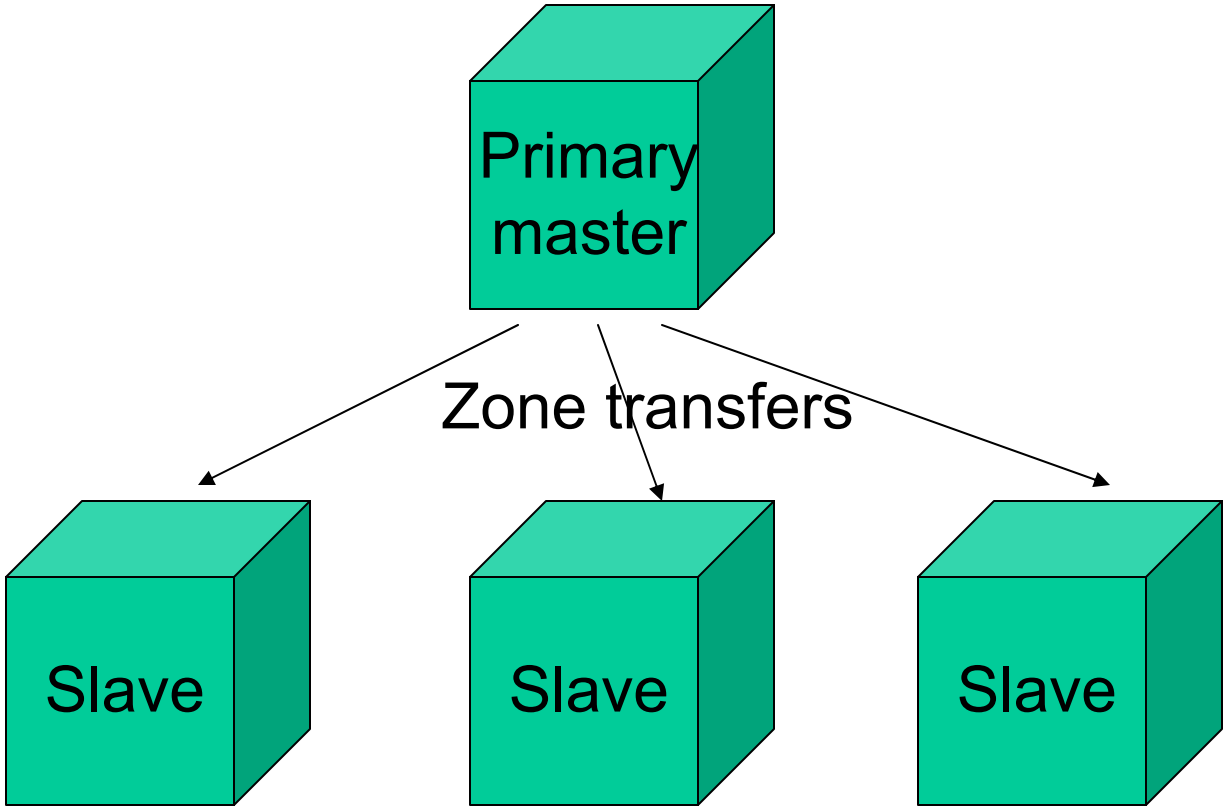

Configuration

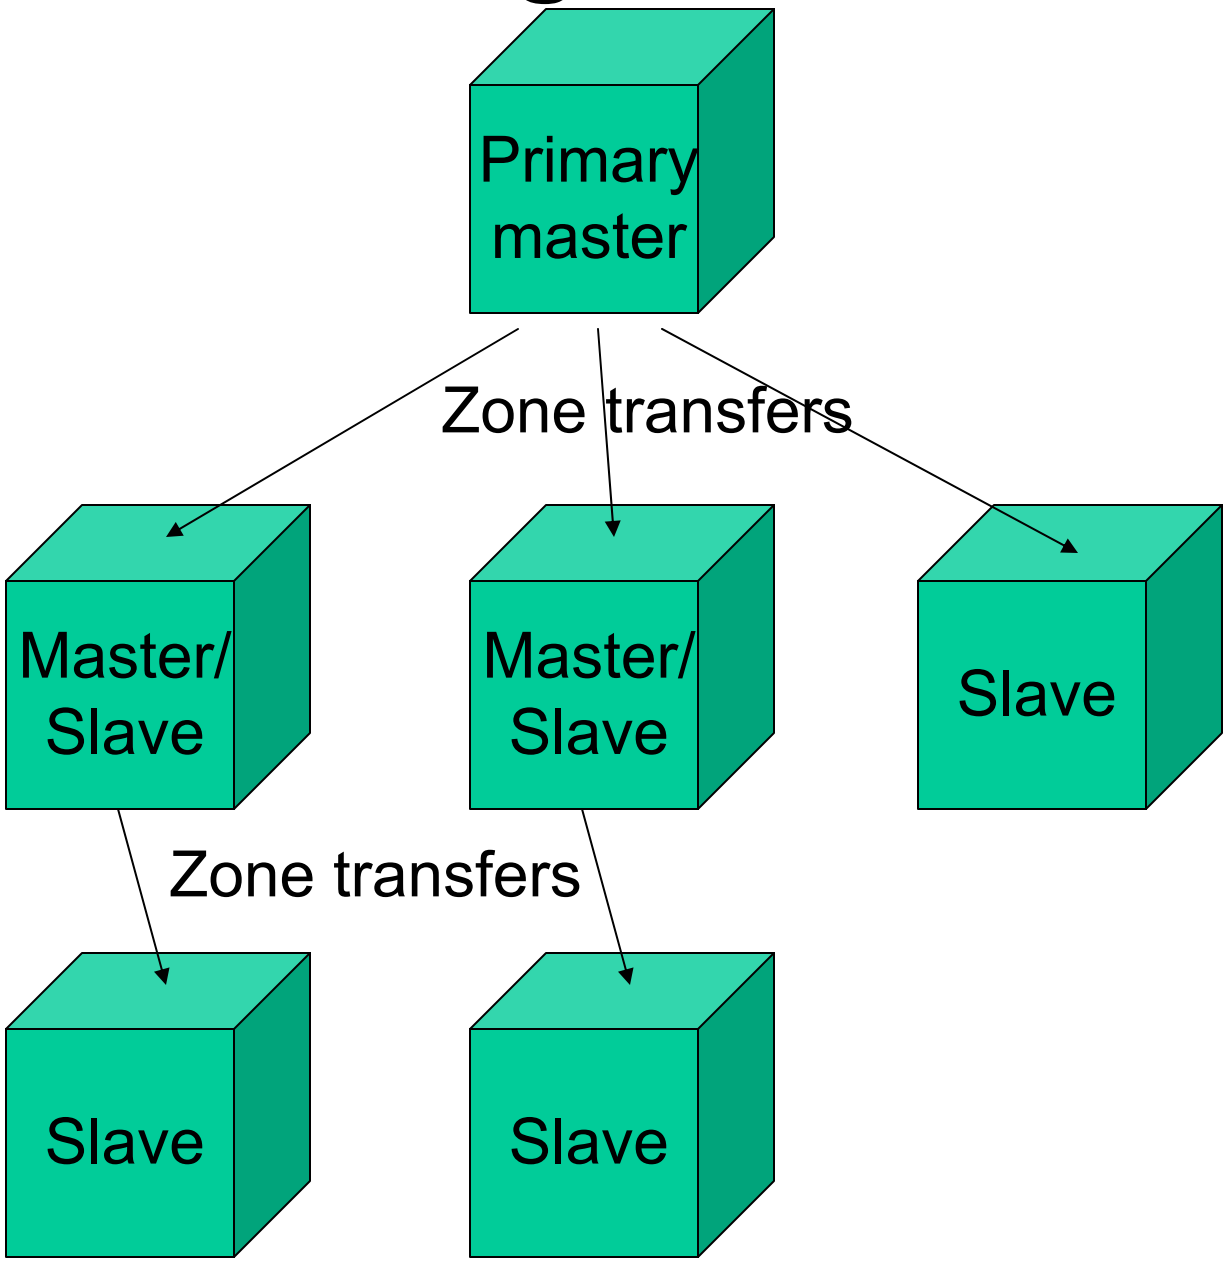

For the Technically Inclined

- ISC provides an “Open Source” reference implementation of the DNS known as BIND
  - Current versions:
    - 4.9.7
    - 8.2.2-P5 (8.2.3 due out soon)
    - 9.0.0 (due out soon)
  - Available from http://www.isc.org/bind.html
- Other freely available DNS implementations
  - DJBDNS, see http://cr.yp.to/djbdns
  - DENTS, see http://www.dents.org
- Commercial DNS implementations available from:
  - Nominum (:-)), Microsoft, Cisco, Lucent, Checkpoint, and many others

Overview

- Introduction
- History
- Name space structure
- Technical details
- Administrative details
- Political details
- Futures
- Summary

Administration of the DNS

- Management of a small scale zones is relatively easy
  - Won't be addressed here
  - See "DNS & BIND, 3rd Edition" by Paul Albitz & Cricket Liu, O'Reilly & Assoc.
- Management of large scale (country sized) zones is a bit more of an issue
  - Services you should provide
  - Traps to watch out for

What Does It Mean to Provide a Country TLD Service

- RFC 920 first documented the concept of country code domains
  - Rumor is ccTLDs were an afterthought
    - Original idea was to have, e.g., all the world's military organizations under .MIL
- RFC 920 was issued in 1984
  - No ccTLD administrators yet existed
  - The IANA initially delegated the ccTLD to anyone who asked
  - Soon, the policy was revised to require the administrative contact in-country
  - The IANA imposed no requirements on how the TLD was to be administered

Providing ccTLD Services

- With international proliferation of the Internet came increased demand for domains from ccTLD
  - The IANA allocated ccTLDs to:
    - Universities
    - Commercial entities
    - Individuals
- Until around 1995 or so, governments ignored the Internet
  - The world was moving to OSI
    - Or so they thought…

Providing ccTLD Services (cont'd)

- As the Internet became more popular/important, national Network Information Centers began to sprout up
  - Entities interested in a domain from a particular ccTLD contacted these national NICs
  - In some cases national NICs provided services equivalent to InterNIC
  - In other cases, the national NIC didn't
- There is no requirement (to date) to provide any services other than name allocation
  - And to insure duplicate names aren't allocated

Useful Services (cont'd)

- Domain name allocation
  - Avoidance of duplicates
- Nameserver operation for the domain
  - Secondary for sub-domains of the domain
- Domain name registration
  - Making the registration database available
- Providing training and information on Internet related issues
- Providing a forum for Internet development within a country

Recommendations for a NIC

- Technical competence
  - Always a good thing
- Good connectivity
  - Required if you'll be running the nameserver primary
- No discrimination
  - Or rather, equal discrimination
    - Fair and equitable policies for all applicants
- Documented policies
  - Including the appeals process

Obtaining the Delegation of a ccTLD

- As of now, only KP (N. Korea) and EH (W. Sahara) are not delegated
- However, increasingly ccTLDs are being transferred between organizations within a country
- To obtain the delegation, you will need:
  - Support of the majority of the local Internet community
    - Or, be a governmental agency if the delegation is not already made to a governmental organization
      - "The desires of the government of the country with regard to delegation of a ccTLD are taken very seriously"
  - Well connected nameservers
    - Albeit, not necessarily in-country
  - Demonstrated technical competence
  - Agreement to comply with RFC 1591

Support of the Local Internet Community

- Everybody should work together
  - Re-delegation will not occur if there is more than one supported contender for the domain name
- No strong objection should exist against the proposed NIC
  - If a strong objection exists, requestor(s) will be told to work it out and come back when they have
- If mistrust exists, TLD may not be re-delegated until a consensus is reached

Well Connected Nameservers

- Required to provide nameservice for the allocated subdomains to the rest of the Internet
- Secondary servers required at other locations
  - Definitely on different networks
  - Preferably in widely geographically dispersed locations
  - Provided free of charge by RIPE-NCC, APNIC, and others
- Bandwidth to nameserver need not be large
  - But must be stable

Proof of Technical Competence

- NIC's nameservers control all domains under it
- Unstable service could result in loss of reachability for all domain owners
  - Unless people have memorized IP addresses
- Database and administrative services must be provided appropriately
  - Backups of critical databases
  - 24x7 systems support pretty much expected
  - 24x7 user support would be nice
    - But extremely rare
- The IANA won't give a test
  - But will pull a delegation if badness happens

RFC 1591 Compliance

- RFC 881 (revised several times, 1591 being the last) set forth the basic principles of name delegation:
  - Technical and administrative contacts must exist and must act as the manager for the domain
  - The designated managers are trustees and serve the community the domain represents
  - The manager must be equitable to all groups
  - Significantly interested parties must all agree that the manager is the appropriate body to administer the domain
  - The manager must do a good job
  - Any transfer of responsibility must be coordinated
- NIC operators will have to indicate compliance with these requirements

Application for a ccTLD

Obtain the template found at

- http://www.iana.org/cctld-template.txt

Fill it in

- It is amazingly similar to the InterNIC domain request template

Send the filled in template to iana@iana.org

Wait…

The IANA will

- Verify the information on the form
- Verify the appropriateness of the request
- Update the InterNIC database appropriate

or

- Send mail back explaining why they won't be processing the request

Operational Requirements

- It is very important to have well formulated and **published** policies
  - Fair and equitable to potential customers
  - Not discriminatory
- It is even more important that the policies are adhered to
- It is helpful to have the policies translated to English and available on the Web
  - People prefer copying to re-inventing the wheel
  - Allows for similar policies over the Internet

Policies (Overview)

- Policies should be defined for the following areas:
  - Structure of the namespace
  - Eligibility for domains under the TLD
  - TLD and subdomain ownership
  - Allocation model
  - Charging and billing
  - Domain name disputes
  - User participation
  - Documentation
  - Appeals process

Structuring the Namespace

- Generally it's possible to just assign any domain under the TLD
  - Flat namespace
- Not generally a good idea
  - Doesn't scale well
  - Can lead to confusion
  - May result in administrative nightmares later
- Structured namespaces are generally better thought of
- Restructuring an existing namespace must provide for a period of transition!

Structuring the Namespace (cont'd)

- If a hierarchical namespace is chosen
  - NIC chooses the SLDs
    - Customers obtain 3LDs
- It is very important to document the circumstances under which a new SLD will be created
  - Community input is generally a good thing
- Policies for determining which SLD should be used should be objectively verifiable
  - e.g., a business license for commercial domains, organization charter for non-profit domains, etc.

Eligibility

- Who is allowed to apply for a domain under the TLD?
  - Only registered entities?
  - Only local businesses, organizations, etc?
  - Anyone in the world?
- What is the policy for multiple domain names by a single organization?
  - If restricted, what is the unit of organization?
    - Subsidiaries? Branches? Offices?

Application Procedures

- Should be explicit and easily understood
- Should be available online
  - WWW, FTP, e-mail
  - Nice to have automated forms with ticketing systems to track requests
- Quality of service parameters should be documented
  - Turnaround time, etc.
- Additional documentation requirement should be spelled out prior to application
  - Minimize the number of query/response cycles

Ownership

- Who owns the delegated domain?
  - The NIC?
  - The organization requesting the domain?
  - The end user of the domain?
- Which rights to the domain are conferred?
  - Full rights?
  - Only the use of the domain for actual activities?
    - As opposed to resale?
- What constitutes (il)legal use of a the domain name?

Allocation of Domains

- Model:
  - Outright sale?
  - Rent/Lease for a period of time?
  - Membership model?
    - E.g., use of the domain name only allowed while owner is a member of the NIC?
- Renewal period?
  - How often
    - Once-and-for-all?
    - Yearly?
- Transfer of domain names
- When can a delegation be revoked?
- When will a domain be re-delegated?
  - What is the NIC's hold time

Charging and Billing

- NICs will incur substantive costs
  - How much will depend on the level and quantity of services provided
- Fees may be charged for NIC services
  - Strongly recommended
    - Only way to insure long term viability
- What is the fee structure?
- How will money be accepted?
  - Credit cards, checks, cash?
  - If foreign ownership allowed, how is currency exchange handled?
- Policy for partial-term names?

Fee Structure

- Flat or differentiated?
  - First domain cheap, next more expensive?
  - Type of organization (non-profit vs. commercial, etc.) determines fee?
- Fee period
  - One time?
  - Periodic?
    - What period? Yearly, monthly, etc.
- Maintenance fees separated from allocation fees?

Profits

- If fees are charged, they should at least cover costs
- Domain names are now seen as a valuable resource
  - People will question right of the NIC to make money off that resource
- If no competition is allowed, non-profit status might reduce flamage
  - Alternatively, profits can be reinvested into the Internet community
  - Alternatively, people can be told to get a life

Dispute Resolution Policy

- Domain names can conflict with trademarks
- Disputes **will** arise:
  - Domain names must be unique
  - Trademarks need not be
- NICs have been sued in the past
- If a lawsuit is filed, costs will be incurred
  - Expect this and plan for it in budgets
- If a lawsuit is lost, operation of the NIC can be imperiled
  - It would be nice if the NIC is recognized as non-sue-able by the government

Avoid Being Sued

- Not equal to avoiding disputes
- Require customer to stipulate non-infringement
  - Make it their responsibility
- Require indemnification
- Insure delegation does not confer any legal right to the *name* that corresponds to the domain
- Evaluate whether to check for infringements carefully
  - Don’t become an “editor” of content

Dispute Policy

- Clearly define where customers claim infringement
  - If at all possible, not at the NIC
- Define how parties involved will participate in the resolution of the dispute
- Indicate how disputes will be handled
  - Whether the domain will be put “on hold”
    - Whether it will be usable in that state or not
- The NIC should try **really** hard to avoid adjudicating disputes
  - That’s what court systems are for

Overview

- Introduction
- History
- Name space structure
- Technical details
- Administrative details
- Political details
- Futures
- Summary

DNS Politics -- Background

- In 1992, the National Science Foundation re-bid the Network Information Center function
  - The new NIC would be called “InterNIC”
    - Composed of 3 parts
      - Registration services
      - Database services
      - Information services
  - Network Solutions won registration services
  - AT&T won database services
  - General Atomics won Information services
- Transfer from SRI-NIC didn’t go smoothly
  - But after a while, everything worked pretty well…

Mid-Term Review

- In 1995, NSF commissioned a mid-term review for the InterNIC project
  - General Atomics failed
    - Dropped from the InterNIC cooperative agreement
  - AT&T passed, but just barely
    - Told to do more
  - NSI passed with flying colors
    - But it was noticed NSI was struggling under a significant load
    - Increasingly, NSI was being threatened with lawsuits
    - Domain Name Speculators were becoming an issue

Domain Name Growth

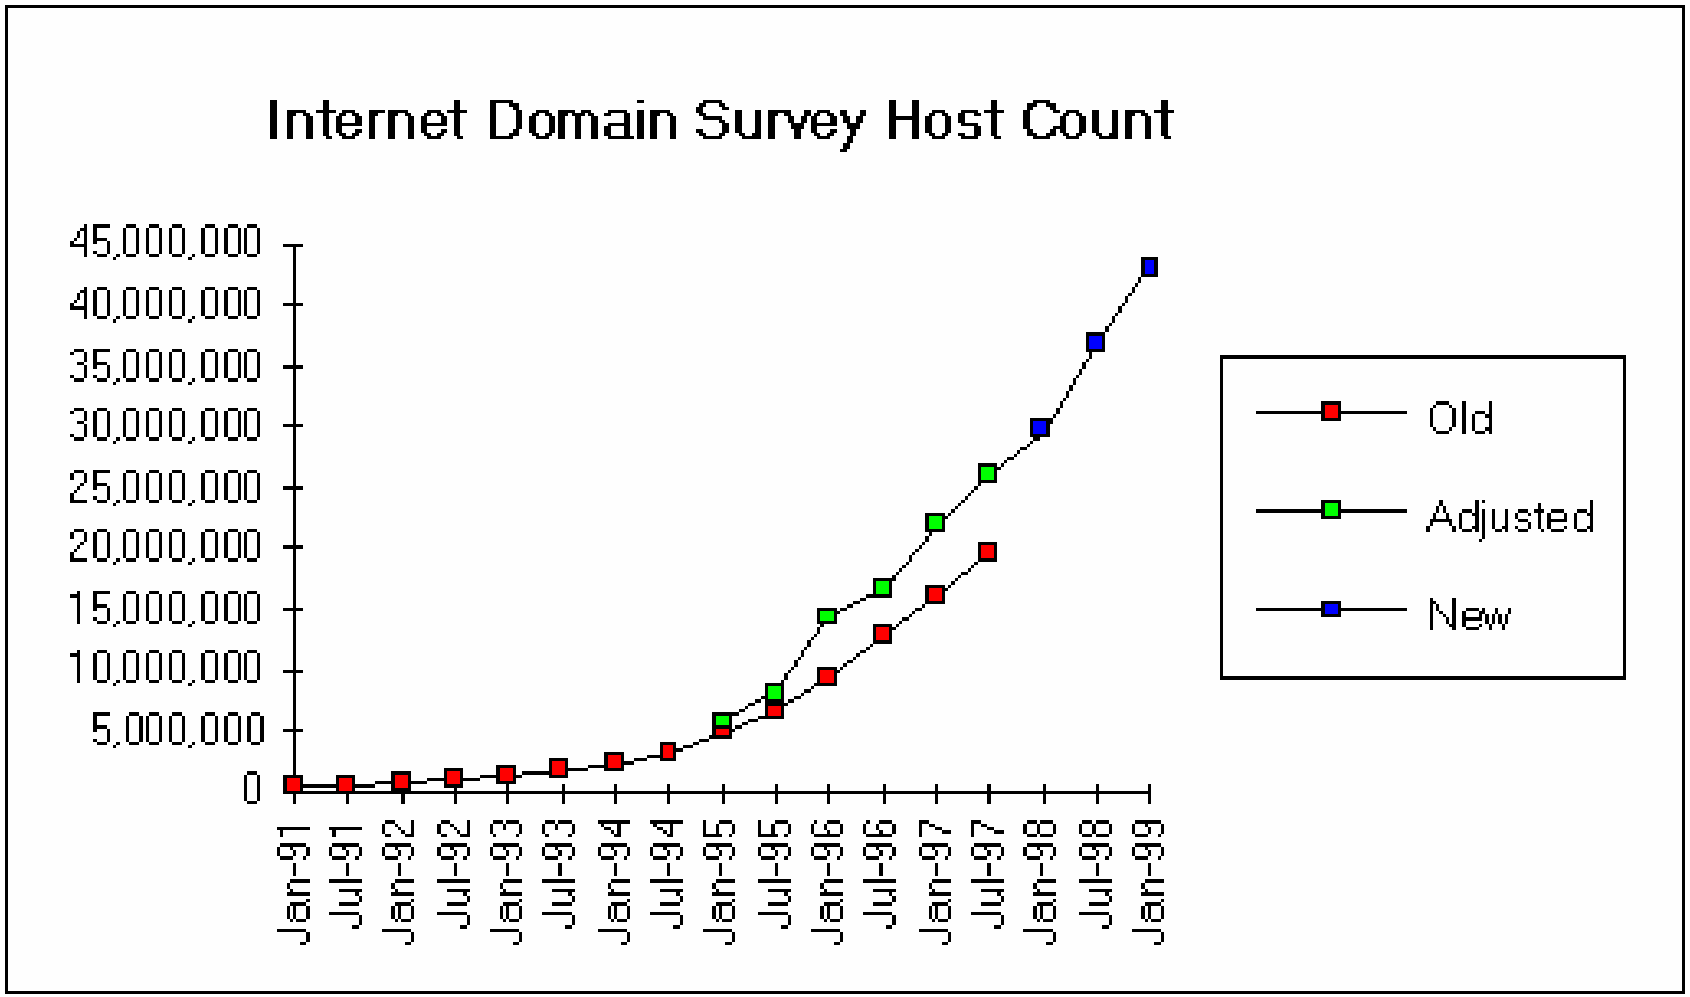

From Network Wizards
http://www.nw.com/zone/hosts.gif

Steps are Taken

- In November, 1995, NSF approves NSI's request to apply a user fee for the allocation of domain names in the .COM zone
  - NSF always intended InterNIC registration services to be self-supporting
    - Funds for RS decreased over time
  - NSI's load related difficulties resulted in NSF paying more money to NSI
- NSF approved a US $50/year domain name registration fee
  - 30% to go to an NSF administered "Internet Infrastructure Fund"

Used Food → Fan

- Many (very vocal) people were outraged
  - "NSI is a government mandated monopoly!"
  - "No competition!"
  - "Infrastructure Fund is a tax!"
- NSI's dispute resolution policy further enrages the masses
  - NSI policy favors trademark holders
- NSI makes a **lot** of money
  - c. 1996, NSI had registered about 2 million domains
- Significant "discussion" ensues

Addressing Concerns

- Early 1996: The YMBK Proposal
  - create small number of new exclusively held TLDs
  - high entrance fee payable to ISOC
- Mid 1996: AlterNIC
  - “Who needs InterNIC anyway?”
  - Point root nameservers elsewhere
- Mid 1996: Open Root Server Coalition
  - Multiple sets of root nameservers
    - Coordinated using out-of-band mechanisms
- May, 1996: Postel Proposal
  - Revised YMBK proposal
    - lower fees, clarified requirements

IAHC

- In Sept. 1996, Postel throws the problem to ISOC
- ISOC formed the “International Ad Hoc Committee”
  - Composed of people nominated by various “stakeholders”
    - IANA, IAB, WIPO, ITU, INTA, ISOC
- in Dec. 1996, IAHC came up with a proposal to create 7 new TLDs
  - .firm, .store, .web, .arts, .rec, .info, .nom
- Key feature: TLDs are a public trust
  - all gTLDs must be shared
  - Creation of the gTLD-Memorandum of Understanding
    - Only signatories to the MoU have input into policies

Counter-proposals

- Shared names were (are) a contentious issue
- Counter-proposals focused on creating new non-shared TLDs
  - NSI (not surprisingly) provided one of these proposals
- IAHC proposal modified in response to the counter-proposals
  - Still no consensus
- Multi-root proposals appealing (decentralizes the DNS), but fundamentally flawed
  - The DNS just doesn't work that way

Resulting Key Concepts

- Competition is the goal
  - How to get there is the question
- Separation of Registration Functions
  - Registry -- entity operating the database containing registration information
  - Registrar -- entity submitting add/delete/change requests to the registry
  - Registrant -- entity requesting the registrar perform add/delete/changes to a domain name registration
- For example:
  - NSI historically has acted as both registry and registrar for .COM, .NET, and .ORG

U. S. Government Steps In

- Jan 30, 1998, US Government issues the “Green Paper”
  - Draft policy document
  - Unilaterally asserts US Government can decide Internet Governance policy
- Proposes
  - to create a new non-profit entity to take over IANA functions
    - Very specific definitions of how it would be created, who would be on the board
  - competition at the registrar level, non-shared registries
  - creation of up to 5 new gTLDs
  - registries should take steps to avoid trademark infringement
  - Internet Infrastructure fund administered by NSF
- Response
  - No significant complaints about US Gov’t assertions of control
  - Much unhappiness about the level of detail
    - Many felt the Green Paper much too “top down”

The White Paper

- Jun 5, 1998, US Government issues the “White Paper”
  - A Statement of Policy
- Revision of the Green Paper, taking into account input received during the comment period
  - Dodges all the hard problems
    - Describes Board of non-profit, but no indication of how to select the board
    - Non-profit to establish criteria for new TLDs, but no hint of what those criteria should be
    - Competitive vs. non-competitive registries left for further study
    - Ask WIPO to propose solution to trademark issues
  - However, this is what people thought they wanted
    - Non-profit should decide as much as possible

ICANN

- Internet Corporation for Assigned Names and Numbers
  - "Newco" (the non-profit) described in the White Paper
- Formed in Oct., 1998
- Composed of
  - Board of
    - 19 directors
    - 9 at-large directors elected by supporting organizations
    - 1 president/CEO (ex officio board member)
  - A secretariat
  - 3 Supporting Organizations
    - Address Supporting Organization
    - Domain Name Supporting Organization
    - Protocol Supporting Organization

DNSO

- The DNSO “will advise the ICANN Board with respect to policy issues relating to the Domain Name System”
- Composed of
  - A Names Council
  - A General Assembly
  - Various Constituency Groups representing specific interests
    - Currently 7 are defined
      - 6 have been formed
    - Constituencies recognized by the ICANN board

DNSO Names Council

- Consists of 3 representatives from each constituency
- Responsible for the consensus building for DNS policies
- Policies can be proposed by the GA or the any one of the Constituencies
- Decisions are made by 2/3's consensus
- NC members nominated by the GA
  - Simple majority of NC members confirms appointment

DNSO Constituencies

- Constituencies are self-organized
  - Recognized by the ICANN board by a majority vote
- Initial Constituencies are:
  - ccTLD registries
  - commercial and business entities
  - gTLD registries
    - Currently only NSI
  - ISP and connectivity providers
  - non-commercial domain name holders
    - The contentious one
  - registrars
  - trademark, other intellectual property and anti-counterfeiting interests

Fun With ICANN

- Many complaints about closed-ness
  - Some decisions made with no indication as to how the decision was reached
  - Closed meetings in which policy decisions were made
- Concerns about how the interim board was chosen
  - By whom? Under what conditions?
- Unhappiness with the decisions made
  - The current board is “interim”
- However…
  - No other game in town
  - ICANN recognized by US Dept. of Commerce as “newco”

More Fun With ICANN

- Governmental Advisory Committee (GAC)
  - Closed committee consisting of governmental appointees
    - Provides input to ICANN
  - Some are concerned about how binding that input is
    - Argument is that it is the GAC that gives ICANN its legitimacy
  - GAC as ejected some designated representatives from meetings
    - Individuals were not governmental employees
      - Worked for US based DNS registrar handling the country's TLD

Even More Fun with ICANN

- Non-commercial Domain Name Holders Constituency
  - Having significant trouble self-organizing
    - ISOC in one camp, ACM in the other
  - Key issues:
    - Exclusion of individuals in other constituencies
    - Definition of “non-commercial”

Individual Constituency

- Proposal has been made for an 8th constituency
  - Created to “provide representation in the DNSO for all Domain Name Owners, who do not wish to be classified as non-commercial, nor wish to be represented by the Business constituency.”
  - Mission is “to be a voice for the Individual DN owners in the DNSO, to work for their interests and to provide representation for them on the Names Council”
- Proposal sent to the ICANN interim board at the Berlin Meeting (May, ‘99)

What Does it All Mean?

- The Domain Name System is fundamentally broken
  - DNS solves a technical problem in a technical way
    - Social/political issues were ignored
  - TCP/IP is designed to be decentralized
  - The DNS requires centralization
    - The root
- This centralization has attracted those interested in controlling the Internet
  - Really amazing amounts of political machinations going on
- ICANN's increasingly Byzantine structure is a result of those machinations
  - Don't expect it to get better anytime soon...

Overview

- Introduction
- History
- Name space structure
- Technical details
- Administrative details
- Political details
- Futures
- Summary

Futures (Technical)

- The DNS protocol suite continues to evolve
  - Future improvements will include
    - stronger DNS security
      - DNS could be the basis for a PKI
    - support for IPv6
      - Should that be necessary
    - Integration with directory service protocols
      - LDAP, in particular
- DNS is being used for things far outside the original scope
  - When all you have is a hammer, everything looks like a nail

Futures (Administrative)

- Policies and procedures for TLD administration will change
  - Continued migration of TLD administration to governmental bodies
  - Clearer definition of operational requirements
  - More/clearer policies and procedures will be required of the TLD administrators
- ccTLDs have strong advantages in the intellectual properties arena
  - There is a non-controversial place to sue
  - Possible migration (duplication) of current non-ccTLD names into the ccTLD spaces

Futures (Political)

- ICANN will stabilize
  - Really is no other game in town…
- Continued political jockeying
  - Becoming a player in the Internet has become a political prize
- Domain Names are seen as big money potentials
  - Many people want their cut
  - Arbitrary fees in the gTLD space will likely push people to the stability of the ccTLDs

Overview

- Introduction
- History
- Name space structure
- Technical details
- Administrative details
- Political details
- Futures
- Summary

Summary

- The Domain Name System provides human-friendly identifiers for Internet users to reference the sites they want to get to
- The DNS is a distributed, global name lookup system composed of
  - A namespace in which all names reside
  - Resolvers which are the clients in name lookups
  - Servers which respond to queries from clients
- The Namespace is hierarchical with a **single** root
  - Required for consistency
  - Provides a handle for the politicians

Summary (cont'd)

- After length "discussion" administration of the DNS has been vested in ICANN
  - As delegated by the US Government
- ICANN still being formed
  - Interim board is making policies
  - Domain Name constituencies being established
  - Battle lines being drawn

"What's in a name? That which we call a rose
by any other name smells as sweet."

William Shakespeare
Romeo and Juliet

Where to Get More Information

- http://www.rfc-editor.org/rfcsearch.html
  - RFC 1032, 1033, 1034, 1035, 1591
- http://www.isc.org/bind.html
- “DNS and BIND, 3rd Edition”, Cricket Liu and Paul Albitz, O’Reilly & Associates
- http://www.iana.org/domain-names.html
- http://www.icann.org
- http://www.dnso.org
- http://www.ntia.doc.gov/ntiahome/domainname/domainhome.htm
- http://www.networksolutions.com