

Unix System Administration

Chris Schenk
Lecture 12 – Tuesday Feb 27

CSCI 4113, Spring 2007

Time

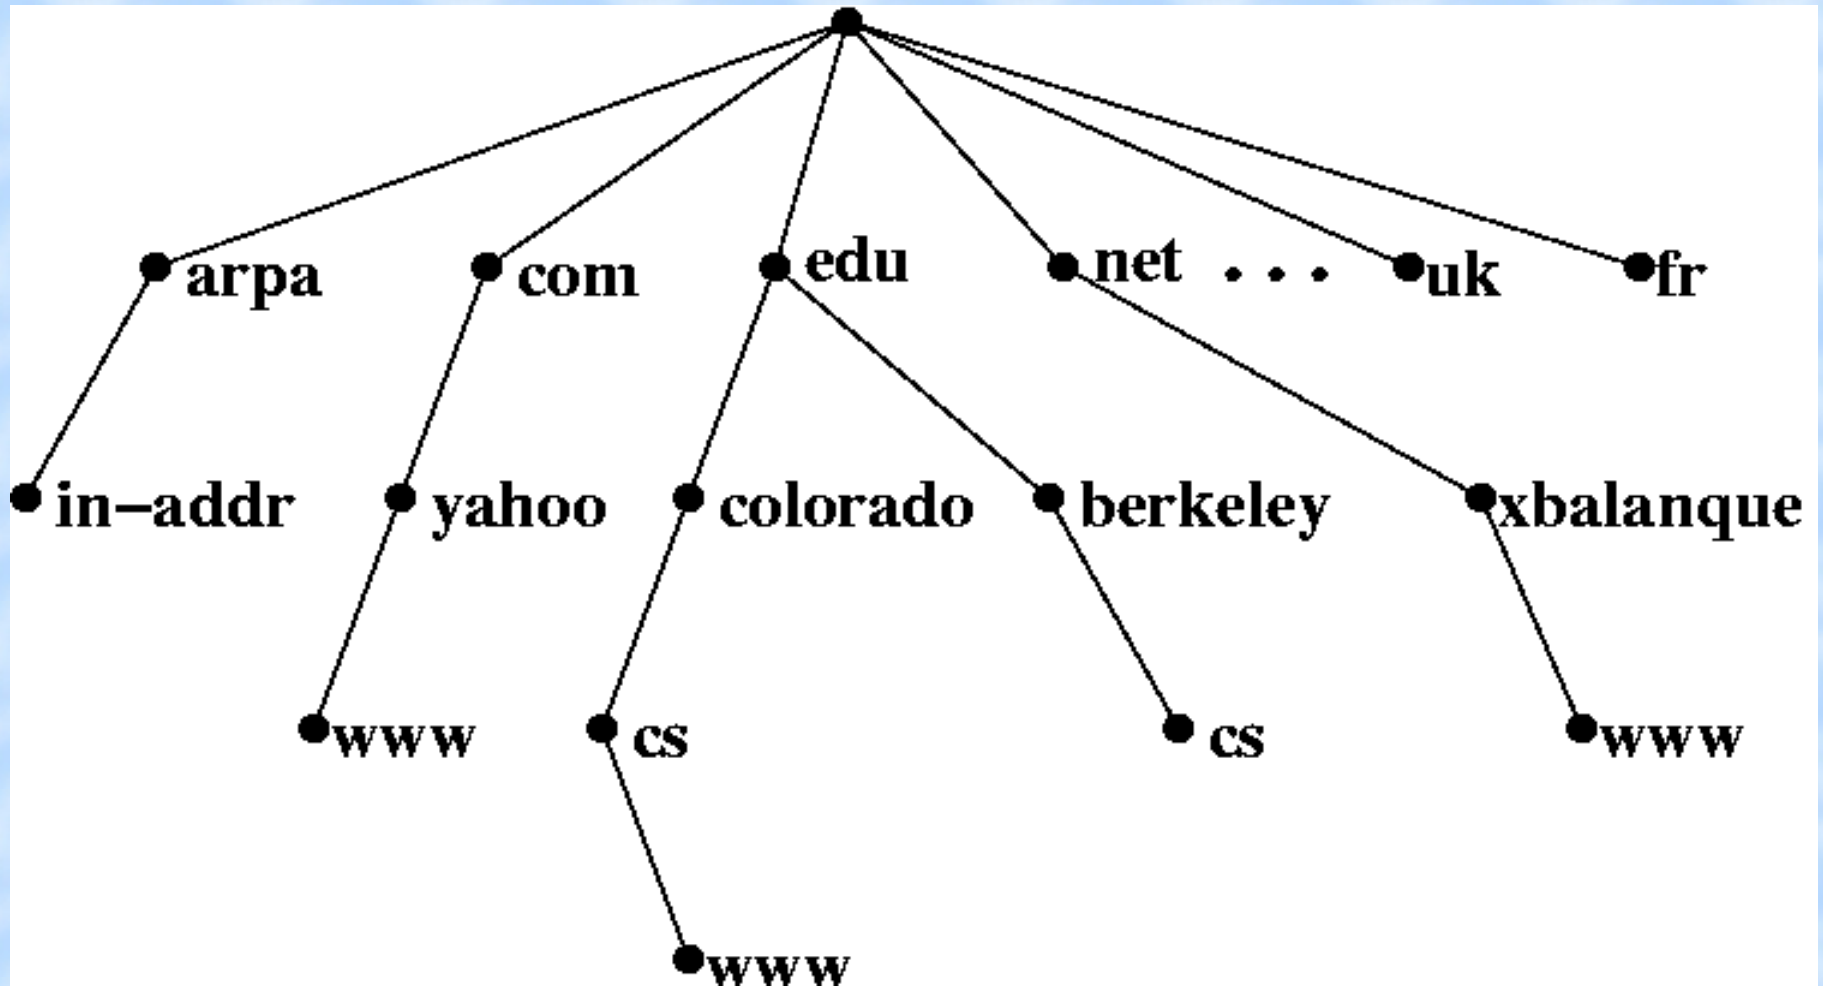
- Marches infinitely?
 - Not with a 32-bit counter!
- Unix time standard is with a 32-bit counter
 - Starting at the “Unix Epoch”
 - Midnight UTC January 1, 1970
 - Counts number of seconds since the epoch
- Another Y2K problem?
 - 03:14:08 UTC Jan 19, 2038 – 32-bit overflow

Time Keeps On Slippin'

- Network Time Protocol (NTP)
 - Allows for synchronization of clocks
 - Necessary for certain things to function
 - sudo, nfs, and many more
- Protocol written to deal with latency
 - Network latency when querying a time server!
 - Also tries to learn the drift of your clock
 - So it can adjust without as many queries

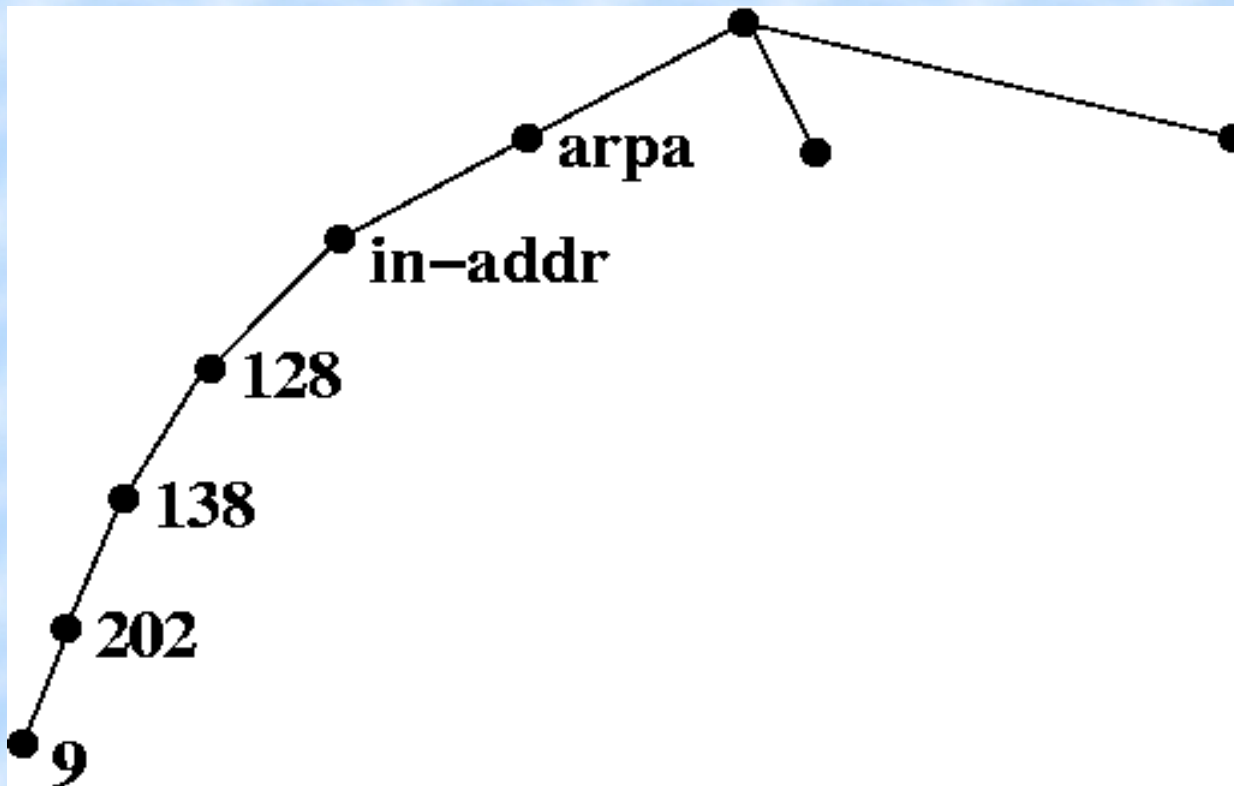
DNS – Domain Name Space

- Distributed database framework



DNS – in-addr.arpa

- This subtree is used for reverse lookups



DNS – Zones and Delegation


- Authority for a zone is delegated by a parent zone
 - The **root** zone delegates **edu** authority to people who manage the **edu** zone
 - The **edu** zone delegates **colorado.edu** authority to the University of Colorado
 - CU then delegates the **cs.colorado.edu** zone to Computer Science
- Delegation of authority means that some parent nameserver has data records that associate a given child zone with some **specific** nameservers

DNS – Zones and Subdomains

- A zone can include its subdomains
 - child zones don't have to be delegated out
- Subdomains are in the parent zone if no **glue** records exist to delegate authority
- A zone is a point of delegation in the DNS tree
 - Don't get this confused!
 - A zone consists of these contiguous parts for which a name server has complete information authority
- Delegation points are marked by one or more NS records

DNS – Zone Delegation Example

- foo.com has multiple domains available
- research.foo.com likes to do their own thing



DNS – Zone Delegation (cont)

- Zone authority usually includes at least two zones
 - Forward and Reverse zones together
- Usually each separate subnet has its own reverse zone
 - This is done even though there is a single forward zone
 - cs.colorado.edu has one forward zone but many reverse zones
- The term 'zone' often refers to both forward and reverse
 - Hardly ever is a forward zone delegated without the reverse or vice versa, but it can happen!

DNS – Queries Revisited

- DNS clients contact nameservers for host resolution
- Local DNS servers may be authoritative for local names and addresses
 - cs.colorado.edu domain is managed locally
- All other queries must be re-transmitted to nameservers that are authoritative for the information requested
 - Exactly how this process works is at the heart of DNS as a robust, distributed database

DNS – Query Resolution Example

- Here's how most standard queries are resolved:
 - I point my browser to www.cs.berkeley.edu
 - A recursive query is sent to my local nameserver (1st in resolv.conf)
 - What is IP address for www.cs.berkeley.edu?
 - If the server doesn't immediately know in the cache, query is sent to the ROOT nameserver
 - Referral received to talk to an EDU server
 - Local server sends query to EDU server
 - Receives referral to berkeley.edu server

DNS – Query Resolution Example (cont)

- Local server queries berkeley.edu server
 - Receives referral to cs.berkeley.edu server
- Finally the query is sent to cs.berkeley.edu
 - And the IP address of www.cs.berkeley.edu is returned
- The local server then sends the response back to the client
 - Client uses IP address to talk to 169.229.60.205 directly

DNS – Resource Records

- There are dozens of resource records
 - Only a few are used to any great extent
- Two records are categorized as 'zone' records
 - SOA “Start of Authority” record – only ONE per zone
 - Indicates the beginning of a zone and should occur first in file
 - Defines values for serial number and expiration times
 - NS “Nameserver” record
 - Defines an authoritative nameserver for a zone
 - Delegates authority to a nameserver for a child zone
 - NS records are the GLUE that binds the distributed database together

DNS – Host Records

- Records that deal mainly with host information
 - A Address record
 - Name->IP address **forward** record
 - PTR Pointer record
 - IP->Name **reverse** record
 - CNAME Canonical Name record
 - Aliases for A records
 - MX Mail eXchange record
 - Specifies a host that will accept mail on behalf of another host

DNS – MX Records

- A single host may have multiple MX records
 - Records for a host make up a priority list
 - The LOWEST number is HIGHEST priority
 - If delivery cannot be made to highest priority, the next highest is attempted
- The idea is highest priority is final destination for email
 - MX backup hosts hold mail to attempt delivery to final host
- MX records used to be optional
 - Many mail configurations won't deliver mail to a host w/out MX

DNS – Security Records

- Records designed to guarantee authenticity of records
 - **KEY** Public key record
 - Associates a public key with a DNS name
 - Same public/private keypair idea
 - **SIG** Signature record
 - Cryptographically sign the resource record
 - **NXT** Next Record
 - Bit misleading
 - Used to authoritatively deny the existence of a record
 - i.e. “host xyz doesn't exist in zone foo”

DNS – Nameservers

- Daemon process that answers DNS queries
 - Come in three basic types: Primary, Secondary, and Caching
- Usually there is one primary server for a zone
 - aka Primary or Master server
 - Primary server loads zone information from files on disk
 - The term 'zone' refers to both forward and reverse zones
- Most sites have one or more secondary servers
 - aka Secondary or Slave server

DNS – Nameservers (cont)

- Secondary servers load zone information from the master
 - All data is transferred at one time via a “zone transfer”
 - Newer implementations of DNS support incremental transfers
- Secondary servers get new data by comparing serial #s
 - Serial numbers found in the SOA record
 - If serial number on slave is less than serial number on master
 - then transfer all new data to slave from master

DNS – Nameservers (cont)

- Third type is the Caching nameserver
 - Caching servers don't serve data for any specific zone
 - Server merely remembers answers for queries it has made
 - Increases network efficiency because repeated queries don't need to be re-sent
 - Cached answer can be used instead
- One nameserver CAN perform all three roles
 - Server can be primary for some zones, secondary for others
 - Can also cache answers for all other queries outside of above zones

DNS – Nameservers (cont)

- Final category of name server is the Stealth server
 - By definition, stealth servers don't have NS records pointing to them
 - NOT compliant with DNS RFC specifications
 - You have to know the server's IP address to use it
 - Stealth servers can be authoritative for a zone
 - Very beneficial for running on a network to increase performance
 - Stealth servers are very common in NAT'ed environments
 - CSEL internal network uses a stealth server

DNS – Zones

- Every zone hosted by an authoritative server
 - Both Primary and Secondary servers authoritative
- Having more than one Primary server is bad
 - Have to manually synchronize data in the zone between servers
- Should have at least one Secondary server for the zone
 - Required by the RFCs to run an official domain
- Frequently the Secondary server will be offsite
 - Often two sites will host Secondaries for each other
 - I'll scratch your back, you scratch mine

DNS – Parent Zones

- Parent zones must delegate to all authoritative servers
 - Primary and Secondary nameservers for a zone should all be referenced with NS records in the Parent zone
 - Many sites only have a select subset of servers referenced by the parent
 - Useful for staying secure by only allowing exposure of certain hosts
 - When in reality there could be many DNS servers for a zone
 - Usually only 2-3 servers are referenced by the parent
 - Again, this is the GLUE that binds the DNS database together

DNS – Forwarding

- Name servers can forward all queries to another server
 - Stealth servers that may not be able to make outgoing queries due to firewalling
 - These servers cache the answer given
 - But they don't learn about any intervening nameservers
- Sites will designate some beefy servers as NON-forwarding servers
 - All other DNS servers forward requests to these ones
 - This builds a large cache of DNS data on the beefy servers

DNS – Network Information

- Servers listen on port 53 both TCP and UDP
 - Port 53 is listed under /etc/services
 - ```
% grep domain /etc/services
domain 53/tcp nameserver
domain 53/udp nameserver
```
- Most queries use UDP
  - Single packet goes to server with a single packet response
- TCP is used for 'zone transfers'
  - Large amount of data to transfer, TCP makes more sense

# DNS – Zone Files

- All Primary zone data is loaded from zone files
  - One zone file per Primary zone
  - File format is standard and is independent of implementation
- Zone files usually maintained by hand
  - Sometimes files are generated by scripts for manageability
  - Straight ASCII files with zone data
- Most entries in zone files are Resource Records (Rrs)
  - Few other control directives and options
    - \$TTL, comments

# DNS – Zone Files (cont)

- First entry must be the \$TTL directive
  - Sets the default Time-To-Live on all data in the zone
    - The amount of time, in seconds, that a remote server will cache the data
  - Syntax is simple
    - \$TTL 7200
  - Number of seconds is specified as an integer
  - \$TTL is not specified in the SOA record
    - Used to be

# DNS – Resource Record Format

- `[name] [ttl] [class] <type> <data>`
  - The name is a domain and is sometimes optional
  - The '@' sign is shorthand for the domain-name of the zone
    - Can be changed with the \$ORIGIN directive
  - When consecutive records refer to the same name
    - Only first record must specify the name
    - Subsequent records **MUST** start with a blank space
    - When names **ARE** specified they **MUST** start at column one
  - Unqualified names (i.e. names without a trailing dot) have the domain appended to the name

# DNS – Resource Record Format (cont)

- Fully qualified names **MUST** have a trailing dot
  - You know when you're missing a trailing dot with names like
  - `csel.cs.colorado.edu.cs.colorado.edu.`
- The TTL record can be set explicitly on a record
  - Although never really seen, always uses the default \$TTL at the top
- The 'class' field defaults to the IN class
  - The 'internet' class, the most common
  - Two other classes are CH from ChaosNet and HS from Hesiod
- Type field specifies the Resource Record type
  - e.g. A, NS, TXT, etc