

# PC & Network Security

CNET – 250 Section 01 CRN: 19234  
Fall 2011

David L. Sylvester, Sr., Assistant Professor



# Chapter 6

## Network Security II

# The Application Layer and DNS

The physical, link, network, and transportation layers provides a basic underlying network infrastructure that allows applications to communicate with each other. It is in the application layer that most of the action of the internet takes place.

## Application-Layer Protocols

There are many application-layer protocols designed to perform a number of important tasks at internet-scale, including the following.

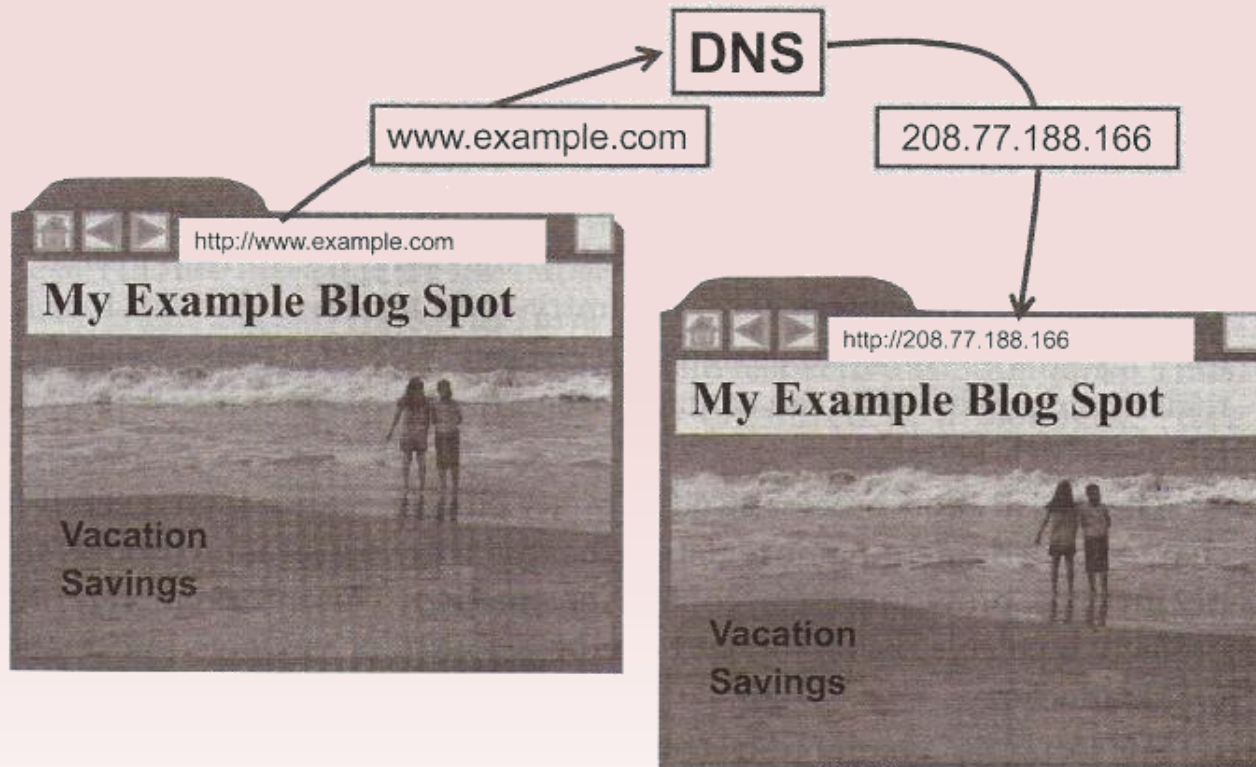
- **Domain Name System (DNS)** – this protocol allows us to use intuitive domain names to refer to internet hosts rather than using IP addresses. Most application programs and other application-layer services rely on DNS.
- **Hypertext Transfer Protocol (HTTP)** – this protocol is used to browse the web.

- **SSL/TSL** – this protocol is used for secure, encrypted browsing (i.e. HTTPS)
- **IMAP/POP/SMTP** – these protocols make internet email possible.
- **File Transfer protocol (FTP)** – This is an old, but still used protocol that provides a simple interface for uploading and downloading files. It does not encrypt data during transfer.
- **SOAP** – This more recent protocol is used for exchanging structured data as a part of the web services paradigm.
- **Telnet** – this is a remote access protocol.
- **SSH** – this is a more recent secure remote access and administrative protocol.

Each application-layer protocol comes with its own security considerations.

# The Domain Name System (DNS)

DNS is a protocol that sits “behind the scenes” for every web browser and is responsible for resolving domain names, such as [www.example.com](http://www.example.com), to IP addresses [208.77.188.166](http://208.77.188.166).



The DNS protocol performs a lookup for domain name [www.example.com](http://www.example.com) to find the IP address associated with this domain.

It is hard to imagine surfing the net without DNS. For instance, would the internet still be popular if we would tell our friends about the video we just watched on 74.125.127.100?

Domain names are arranged in a hierarchy that can be read by examining a domain name from right to left.

Ex: [www.example.com](http://www.example.com) has a top-level domain (TLD) of **com**, with **example.com** being a subdomain of **com**, and **www.example.com** being a subdomain of **example.com**.

The domain name form a rooted tree, where each node corresponds to a its subdomain. The root is the empty domain name and the children of the root are associated with top-level domains.

There are two primary types of top-level domains in use today.

1. Generic top-level domains, such as (**.com**, **.net**, **.edu**, and **.org**)
2. Country-code top-level domains, such as (**.au**, **.de**, **.it**, **.pt**)  
(**Australia**, **Germany**, **Italy**, **Portugal**) respectively

# Generic Top-Level Domains

Name	Entity	Notes
<a href="#">.aero</a>	air-transport industry	Must verify eligibility for registration; only those in various categories of air-travel-related entities may register.
<a href="#">.asia</a> <a href="#">.biz</a>	Asia-Pacific region business	This is a TLD for companies, organizations, and individuals based in the region of Asia, Australia, and the Pacific. This is an open TLD; any person or entity is permitted to register; however, registrations may be challenged later if they are not held by commercial entities in accordance with the domain's charter. This TLD was created to provide relief for the wildly-popular .com TLD.
<a href="#">.cat</a> <a href="#">.com</a>	Catalan commercial	This is a TLD for Web sites in the <a href="#">Catalan language</a> or related to Catalan culture. This is an open TLD; any person or entity is permitted to register. Though originally intended for use by for-profit business entities, for a number of reasons it became the "main" TLD for domain names and is currently used by all types of entities including nonprofits, schools and private individuals. Domain name registrations may be challenged if the holder cannot prove an outside relation justifying reservation of the name, to prevent " <a href="#">squatting</a> ".
<a href="#">.coop</a> <a href="#">.edu</a>	cooperatives educational	The .coop TLD is limited to cooperatives as defined by the <a href="#">Rochdale Principles</a> . The .edu TLD is limited to specific educational institutions such as, but not limited to, primary schools, middle schools, secondary schools, colleges, and universities. In the US, its usability was limited in 2001 to post-secondary institutions accredited by an agency on the list of <a href="#">nationally recognized accrediting agencies</a> maintained by the <a href="#">United States Department of Education</a> . This domain is therefore almost exclusively used by US colleges and universities. Some institutions that do not meet the current registration criteria have <a href="#">grandfathered</a> domain names.
<a href="#">.gov</a> <a href="#">.info</a> <a href="#">.int</a>	governmental information international organizations	The .gov TLD is limited to US governmental entities and agencies. This is an open TLD; any person or entity is permitted to register. The .int TLD is strictly limited to organizations, offices, and programs which are endorsed by a treaty between two or more nations.
<a href="#">.jobs</a>	companies	The .jobs TLD is designed to be added after the names of established companies with jobs to advertise. At this time, owners of a "company.jobs" domain are not permitted to post jobs of third party employers.
<a href="#">.mil</a> <a href="#">.mobi</a> <a href="#">.museu</a>	<a href="#">US military</a> mobile devices museums	The .mil TLD is limited to use by the US military. Must be used for mobile-compatible sites in accordance with standards. Must be verified as a legitimate museum.
<a href="#">m</a> <a href="#">.name</a>	individuals, by name	This is an open TLD; any person or entity is permitted to register; however, registrations may be challenged later if they are not by individuals (or the owners of fictional characters) in accordance with the domain's charter.
<a href="#">.net</a>	network	This is an open TLD; any person or entity is permitted to register. Originally intended for use by domains pointing to a distributed network of computers, or "umbrella" sites that act as the portal to a set of smaller websites.
<a href="#">.org</a>	organization	This is an open TLD; any person or entity is permitted to register. Originally intended for use by non-profit organizations, and still primarily used by same.
<a href="#">.pro</a>	professions	Currently, .pro is reserved for licensed or certified lawyers, accountants, physicians and engineers in France, Canada, NL, UK and the US. A professional seeking to register a .pro domain must provide their registrar with the appropriate credentials.
<a href="#">.tel</a>	Internet communication services	A contact directory housing all types of contact information directly in the Domain Name System.
<a href="#">.travel</a>	travel and tourism industry related sites	Must be verified as a legitimate travel-related entity.
<a href="#">.xxx</a>	adult entertainment	For sites providing sexually-explicit content, such as pornography.

Domain names are registered and assigned by domain-name registrars, which are organizations accredited by the Internet Corporation for Assigned Names and Numbers (ICANN). This is the same group responsible for allocating IP address space or a country-code top-level domain that has been granted authority to designated registrars. Web site owners wishing to register a domain name can contact a domain-name registrar to reserve the name on their behalf.



The screenshot shows the homepage of Active Domain, a domain registrar. The header features the Active Domain logo, the website URL www.active-domain.com, and a McAfee Secure badge. The main navigation bar includes links for HOME, REGISTER DOMAIN, RENEW DOMAIN, TRANSFER DOMAIN, CUSTOMER LOGIN, and FAQs. The central content area is divided into three sections: a Risk-Free Guarantee, a domain registration offer, and a domain availability checker. The Risk-Free Guarantee section states that if there is a typo error or a change of mind after registration, a 72-hour Grace Period is provided for a refund or change. The registration offer highlights a price of \$2.85 per year. The availability checker includes a search box and a 'Search' button. Below this, a 'Why Active Domain is Your Choice!' section lists various free services such as no setup fees, email forwarding, unlimited URL redirections, and 60 days of free trial web hosting. The bottom left corner features a 'Free 60 Days Web Hosting' banner and a 'Learning About Domain Names' button. The bottom right corner includes a photo of a smiling woman.

**AD** www.active-domain.com  
Buy Domain Registration & Register Domain Name Service

McAfee SECURE  
TESTED 30-OCT  
CUSTOMER SUPPORT

HOME :: REGISTER DOMAIN :: RENEW DOMAIN :: TRANSFER DOMAIN :: CUSTOMER LOGIN :: FAQs

**RISK-FREE GUARANTEE**  
If there is a typo error or you change your mind after registering a domain name with us, you have a 72-hour Grace Period to cancel your domain purchase order, get a refund or change to a different domain name.  
\* See [Terms & Conditions here](#)

**REGISTER**  
DOMAIN NAME FOR AS LOW AS  
**\$2.85** /YR\*  
[Buy Domain Now](#)

**CHECK DOMAIN NAME AVAILABILITY**  
eg. mydomain.com    
Enter in your desired domain name into the box above to check its availability before registering.

**FREE 60 DAYS WEB HOSTING**  
WITH EACH DOMAIN ORDER  
All prices quoted in US Dollar  
  


**WHY ACTIVE DOMAIN IS YOUR CHOICE !**  
ENJOY GREAT SAVINGS with our FREE services for each internet domain name registration!

- FREE No setup fee, sales tax or hidden cost
- FREE Email forwarding addresses
- FREE Unlimited URL redirections
- FREE URL frame / cloaking
- FREE Sub-domains or Host Names
- FREE Parking on our DNS Server
- FREE Transfer of Domain Name ownership
- FREE 60 Days Free Trial Web Hosting
- Wide range of Domain Extensions on offer: .com, .net, .org, .info, .cc, .ws, .biz, .us, .uk  
...and lots more! [See details](#)

**LEARNING ABOUT DOMAIN NAMES**



# Availability of the domain name **david** .



The screenshot shows the Active Domain website interface. At the top, there is a navigation bar with the logo and the text "www.active-domain.com". Below this is a green navigation bar with links for "HOME", "REGISTER DOMAIN", "RENEW DOMAIN", "TRANSFER DOMAIN", "CUSTOMER LOGIN", and "FAQs". The main content area is titled "6 Simple Steps to Register Your Domain Name >>". Below this is a progress bar with six steps, where "Step 2" is currently selected. The "Step 2" section is titled "Step 2> Search Results" and includes a shopping cart summary: "Items: 0 | Total: \$0.00 | View shopping cart". The main instruction is "Please select the domains you wish to purchase and add them to your shopping basket:". Below this is a table of search results for the domain "david":

<input type="checkbox"/>	david.com	Not Available
<input type="checkbox"/>	david.net	Not Available
<input type="checkbox"/>	david.org	Not Available
<input type="checkbox"/>	david.info	Not Available
<input type="checkbox"/>	david.biz	Not Available

At the bottom of the search results section, there is a link: "Click here if you wish to search for more domains."

The registration process is pretty simple. Other than a small fee charged by a domain-name registrar, the rest of the registration process simply involves providing some contact information. This information is often publicly available and can be a source of valuable information for an attacker.

The system utility **whois** can be used to retrieve the contact information of the owner of that domain, which might be used to initiate a social engineering attack.



Domain Name: MYBRCC.EDU

**Registrant:**

Baton Rouge Community College  
201 Community College Drive  
Baton Rouge, LA 70806  
UNITED STATES

**Administrative Contact:**

Ronald Solomon  
Chief Information Officer  
Baton Rouge Community College  
201 Community College Drive  
Baton Rouge, LA 70806  
UNITED STATES

 (225) 216-8267   
solomonr@mybrcc.edu

**Technical Contact:**

Lloyd Allen  
IT Security Officer  
Baton Rouge Community College  
201 Community College Drive  
Baton Rouge, LA 70806  
UNITED STATES

 (225) 216-8509   
allenl@mybrcc.edu

**Name Servers:**

NS1.MYBRCC.EDU 64.66.70.68  
NS2.MYBRCC.EDU 64.66.70.67

Domain record activated: 12-Jan-2005

Domain record last updated: 14-Oct-2011

Domain expires: 31-Jul-2012

# whois information on the internet



Online Schools | 4-Year Colleges | Career Training | Community Colleges | Graduate Schools | College Quick Search

## Baton Rouge Community College

201 Community College Drive, Baton Rouge, LA 70806



Need a two-year degree or transfer in Louisiana? Baton Rouge Community College (BRCC) is a public, two-year college that offers just that: transferable associate's degrees and technical training to students in the Baton Rouge area.

Located near downtown Baton Rouge, BRCC is designed to be accessible to eastern Louisiana students, and an affordable provider of quality education.

Some of the programs that BRCC offers include:

- Associate's degrees in Accounting, Business, Construction Management, Computer Science, General Sciences, and Nursing
- Associate transfer degrees in liberal arts and general science, transferable to any Louisiana public university
- Certificate courses in Business, Emergency Management, Customer Service, General Studies, and Diagnostic Medical Sonography
- Corporate training programs
- Continuing education classes
- More than 300 online classes throughout the curriculum

BRCC strives to accommodate all... [Read More](#)

### QUICK FACTS

**Location:** Southeast  
**Setting:** Mid-size City Setting  
**Type:** Public  
**Size:** Large (5,000 to 10,000 Undergrad)  
**Full Time Students:** 54%  
**Athletic Programs:** Unavailable

### Get more information on topics like:

- Admissions
- Expenses
- Course schedules
- Financial aid (check to see if the school offers it and if you qualify)
- And much more!

[GET MORE INFORMATION](#)

### Get Admissions Info on Baton Rouge Community College

(It only takes a minute!)

Step 1 of 2

Choose an Area of Study:

- Select One - 

Start Date:

- Select One - 

Learning Preference:

- Select One - 

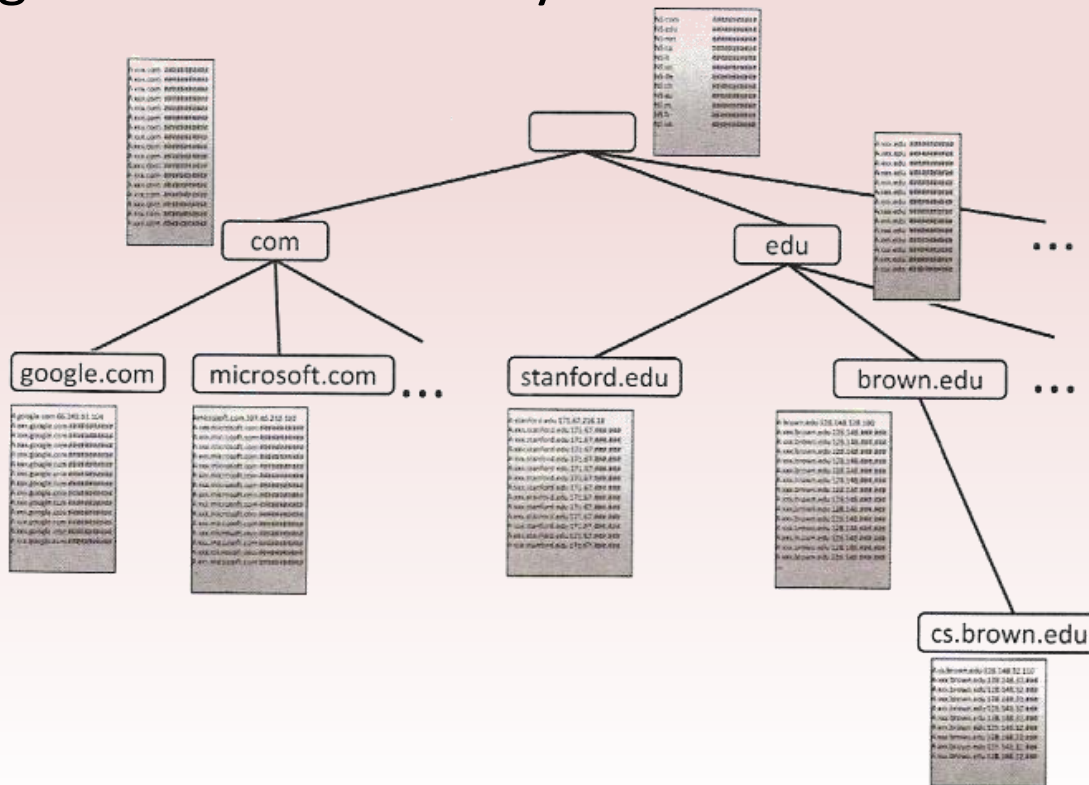
[NEXT STEP](#)

Because of the revenue potential of memorable domain names, a practice known as cybersquatting or domain squatting has become common-place. For example, a person registers a domain name in anticipation of that domain being desirable or important to another organization, with the intent of selling the domain to that organization for what can sometimes be a significant profit.

Some cybersquatters go so far as to post negative remarks or accusations about the target organization on the page to further encourage the target to purchase the domain in defense of its reputation.

How DNS is Organized. At the top of the name-server hierarchy are the root name servers, which are responsible for top-level domains, such as **.com**, **.it**, **.net**, and **.org**. Specifically, the root name servers store the root zone database of records indicating the authoritative name server of each top-level domain. This important database is maintained by ICANN.

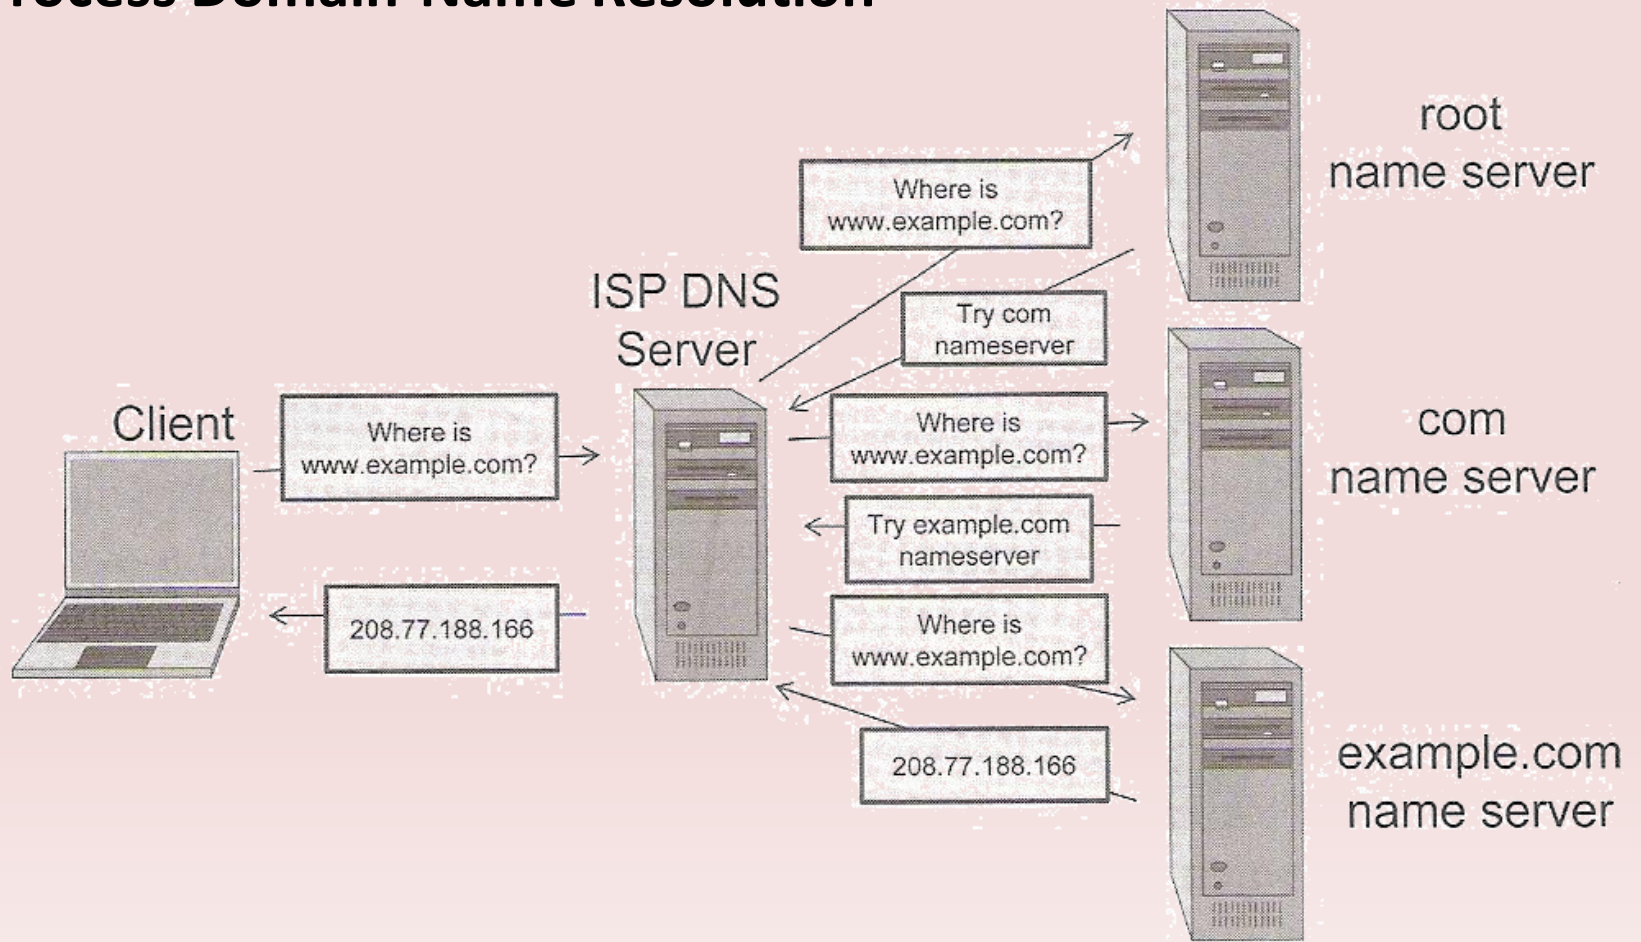
The name servers of each top-level domain are managed by government and commercial organizations. The name servers for the **.com** TLD are managed by VeriSign, a company incorporated in the U.S., while the name servers for the **.it** TLD are managed by the Italian National Research Council, an Italian government organization. In turn, the TLD name servers store records for the authoritative name servers of their respective subdomains. Thus, the authoritative name servers are also organized in a hierarchy.



How DNS Queries Work. When a client machine wishes to resolve a domain name such as [www.example.com](http://www.example.com) to an IP address, it contacts a designated name server assigned to the machine. This designated name server can be a name server of the internet service provider. The designated name server handles the resolution of the domain name and returns the result to the client machine: Ex:

1. The designated name server issues a DNS query to a root name authoritative for the next level of the hierarchy – (replying with the address of the name server for the .com top-level domain name)
2. On querying the next-level server, it would respond with the address of the name server responsible for the next subdomain (example.com)
3. Requests and responses continue until a name server responds with the IP address of the requested domain.
4. The final name server is therefore the authoritative responder for the requested domain name (www.example.com)

# The Process Domain-Name Resolution




The client machine queries a designated name server, such as a name server of its service provider. The designated name server in turn queries a root name server, then a top-level domain name server, and finally the authoritative name server for the requested domain. Once the intermediate name server resolves the domain name, it forwards the answer to the client machine.

DNS Packet Structure. DNS queries and replies are transmitted via a single UDP packet, with TCP being used as a substitute for requests or replies exceeding 512 bytes. The standard UDP packet used in DNS consists of a header, a query part and an answer part.

The header is formatted as follows:

- Header file includes a 16 bit query identifier, also called transaction identifier, which identifies the query and response.



**QR** Flag identifying a query (0) or a response(1).

**Op code** 4-bit field specifying the kind of query: 0 Standard query (QUERY); 1 Inverse query (IQUERY); 2 Server status request (STATUS).

**AA: Authoritative answer** flag. If set in a response, this flag specifies that the responding name server is an authority for the domain name sent in the query.

**TC: Truncation** flag. Set if message was longer than permitted on the physical channel.

**RD: Recursion desired** flag. This bit signals to the name server that recursive resolution is asked for. The bit is copied to the response.

**RA: Recursion available** flag. Indicates whether the name server supports recursive resolution.

**Zero:** 3 bits reserved for future use. Must be zero.

**Rcode:** 4-bit response code. Possible values are: (0 – 5)

The query part is a sequence of “questions”, each consisting of the domain name queried and the type of record requested.

- The query part is a sequence of “questions” (usually just one), each consisting of the domain name queried and the type of record requested. The query ID is selected by the client sending the query and is replicated in the response from the server.



**Length:** A single byte giving the length of the next label.

**Label:** One element of the domain name characters (for example, ibm from ral.ibm.com). The domain name referred to by the question is stored as a series of these variable length labels, each preceded by a 1-byte length.

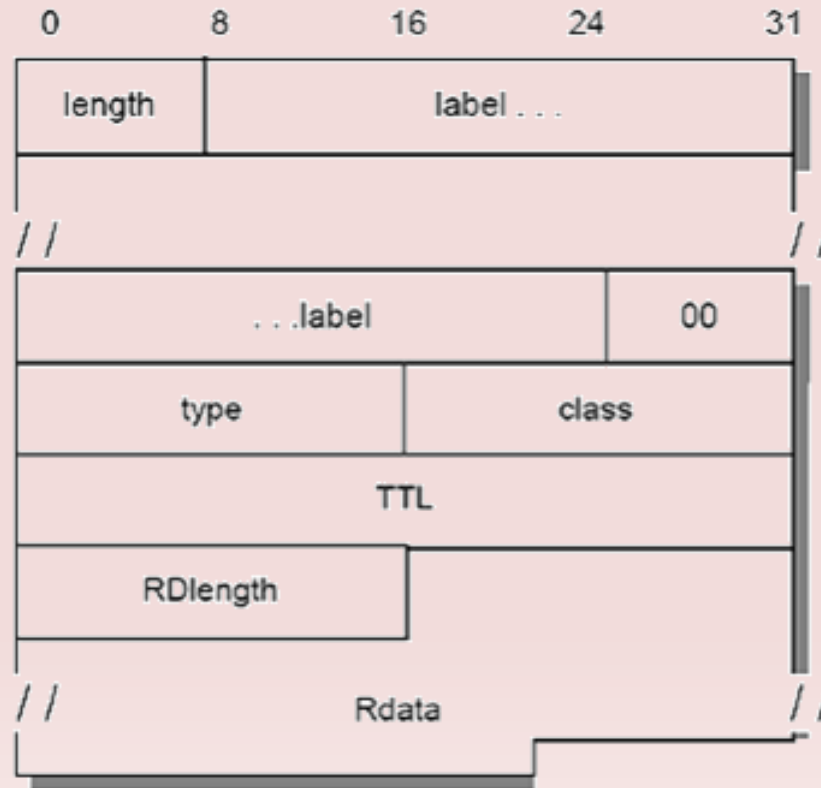
**00:** X'00' indicates the end of the domain name and represents the null label of the root domain.

**Type:** 2 bytes specifying the type of query. It can have any value from the Type field in a resource record.

**Class:** 2 bytes specifying the class of the query. For Internet queries, this will be IN.



The answer part consists of a sequence of DNS records, each consisting of the following fields:



**Name:** The name field is of variable length and contains a full domain name.

**Class:** A field that denotes the broad category that the record applies to, such as IN for internet domain.

**TTL:** A 32-bit time-to-live value in seconds for the record. This defines how long it can be regarded as valid.

**RLength:** A 16-bit length for the Rdata field.

**Rdata:** A variable length string whose interpretation depends on the Type field.

DNS Caching. Since DNS is a central service utilized by billions of machines connected to the internet, without any additional mechanism, DNS would place an incredible burden on high-level name servers, especially the root name servers. In order to reduce DNS traffic and resolve domain name more efficiently, DNS features a caching mechanism that allows both clients and lower-level DNS servers to keep a DNS cache, a table of recently received DNS records. A name server can then use this cache to resolve queries for domain names it has recently answered, rather than consuming the resources of higher-level name servers. This caching system therefore overcomes the problem of massive amounts of traffic directed at root name servers by allowing lower-level name servers to resolve queries.

Caching change how DNS resolution works. Instead of directly querying the root name server each time, the designated name server first checks its cache and returns to the client the requested IP if a record is found. If not, the designated name server queries the root name server and resolves the domain name.

A value known as time-to-live (TTL) determines how long a DNS response record maintains in a DNS cache.

Some operating systems maintain a local DNS cache on the machine. If a valid record is found for the desired domain, then the record is used and no DNS queries are issued. The details of DNS caching depends on the chosen operating system and application.

Ex: Windows features its own DNS cache, while many Linux distributions do not. They choose to query predetermined name servers for each resolution instead.

Several cross-platform browsers, including Firefox, support their own DNS caches. Internet Explorer, which runs on Windows, does not implement this feature because Windows has its own cache.

You can experiment with DNS resolution with the help of several command-line tools. On Windows, nslookup can be used at a command prompt to issue DNS requests. On Linux, users may use either nslookup or dig.

This command is often used to perform a [reverse lookup](#) on an IP address as shown in the below example. The first section specifies the server and address of that server that provided you with the domain name and IP address displayed in the second section.

The reverse lookup, also referred to as reverse resolving and rDNS, Reverse DNS lookup is the process of looking up an IP address to resolve a hostname instead of the other way around. This technique is often used to help diagnose networking related issues or to determine where network data is going.

## Ex: nslookup

**nslookup 204.228.150.3**

Server: ns.computerhope.com

Address: 1.1.1.1

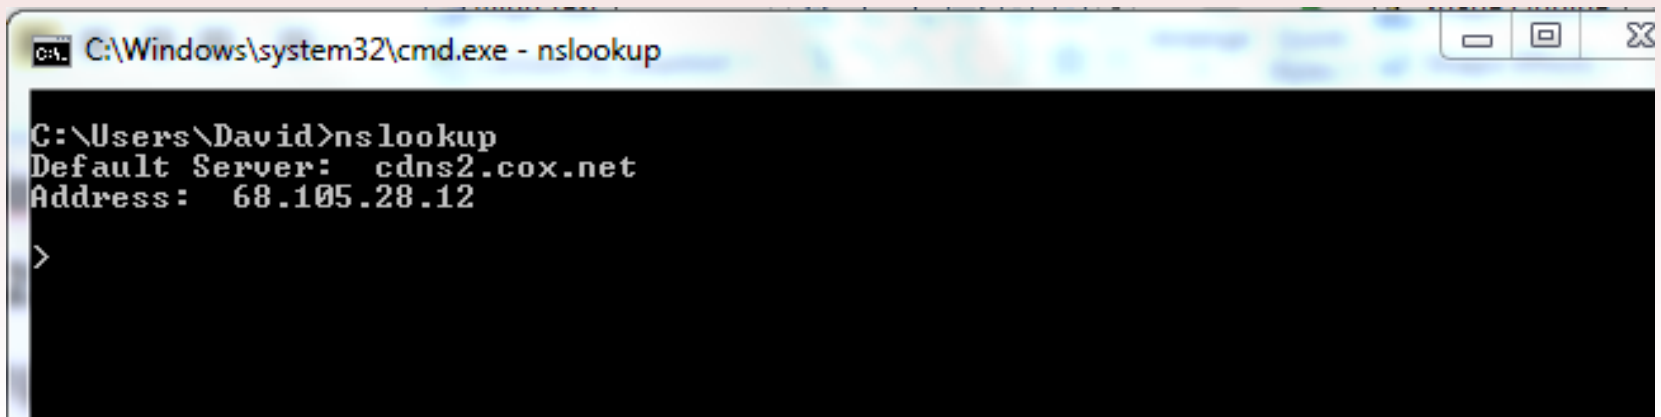
Name: www.computerhope.com

Address: 204.228.150.3

## nslookup

Running nslookup without specifying an IP address or domain name will display your routers server and address. To get out of the > prompt type exit and press enter.

Example.

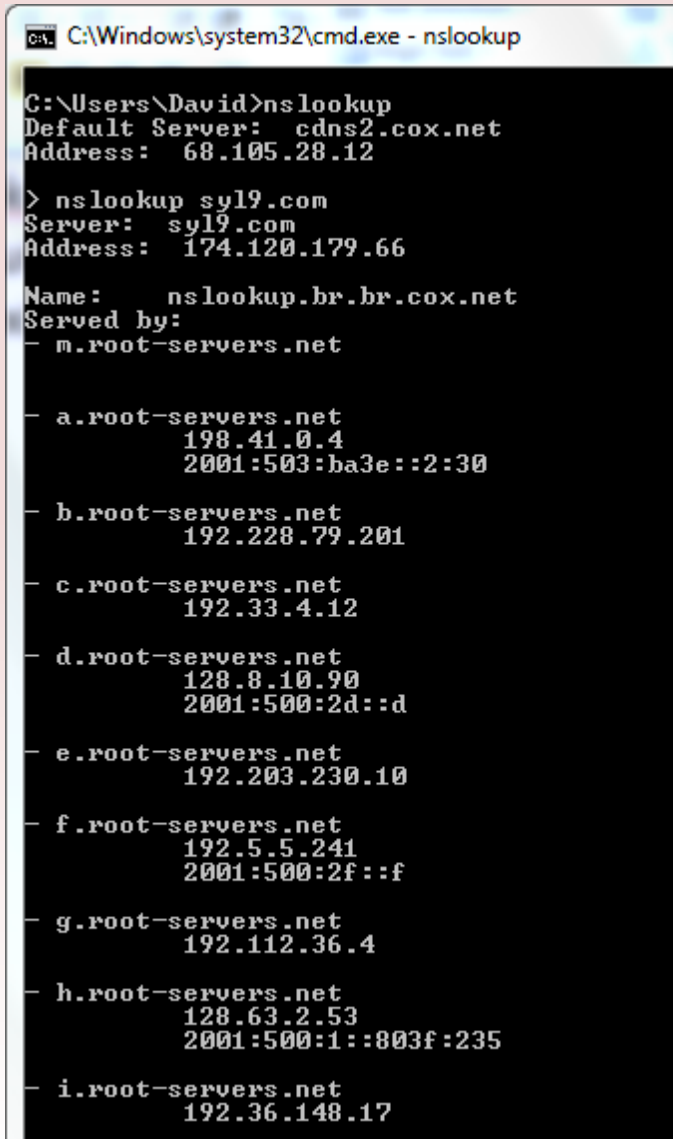


```
C:\Windows\system32\cmd.exe - nslookup

C:\Users\David>nslookup
Default Server:  cdns2.cox.net
Address:  68.105.28.12
>
```

Running nslookup without specifying an IP address or domain name will display your routers server and address. To get out of the > prompt type exit and press enter.

Ex:



```
C:\Windows\system32\cmd.exe - nslookup
C:\Users\David>nslookup
Default Server:  cdns2.cox.net
Address:  68.105.28.12

> nslookup sylv.com
Server:  sylv.com
Address:  174.120.179.66

Name:    nslookup.br.br.cox.net
Served by:
- m.root-servers.net

- a.root-servers.net
  198.41.0.4
  2001:503:ba3e::2:30

- b.root-servers.net
  192.228.79.201

- c.root-servers.net
  192.33.4.12

- d.root-servers.net
  128.8.10.90
  2001:500:2d::d

- e.root-servers.net
  192.203.230.10

- f.root-servers.net
  192.5.5.241
  2001:500:2f::f

- g.root-servers.net
  192.112.36.4

- h.root-servers.net
  128.63.2.53
  2001:500:1::803f:235

- i.root-servers.net
  192.36.148.17
```

# DNS Attacks

By relying on DNS to resolve domain names to IP address, a large degree of trust in the fact that DNS requests are resolved correctly. When we navigate to [www.example.com](http://www.example.com)

Pharming and Phishing. Think about what could happen if DNS were somewhat subverted so that an attacker could control how DNS requests are resolved. Because DNS is so central to how domain names are used to navigate the web, such a subversion would cause the safety of the Web to be compromised. An attacker could cause requests for web sites to resolve to false IP addresses on his own malicious server, leading the victim to view or download undesired content, such as malware.

One of the main uses of pharming is to resolve a domain name to a web site that appears identical to the requested site, but is instead designed for a malicious intent. Such an attack is known as phishing, and it can be used to try to grab usernames and passwords, credit card numbers and other personal information.

Victims of a combined pharming and phishing attack would have no way of distinguishing between the fake and real sites, since all of the information conveyed by the browser indicates that they are visiting a trusted web site.

Example:

Email relies on specialized DNS entries known as MX records, so another possible pharming attack allows an attacker to redirect mail intended for certain domains to a malicious server that steals information. Given that many online services allow password recovery through email, this could be a means of performing identity theft.

An MX (Mail Exchange) record is a type of resource record in the DNS that specifies a mail server responsible for accepting email messages on behalf of a recipient's domain.

*Randomization of transaction IDs is one deterrent for DNS cache poisoning.*



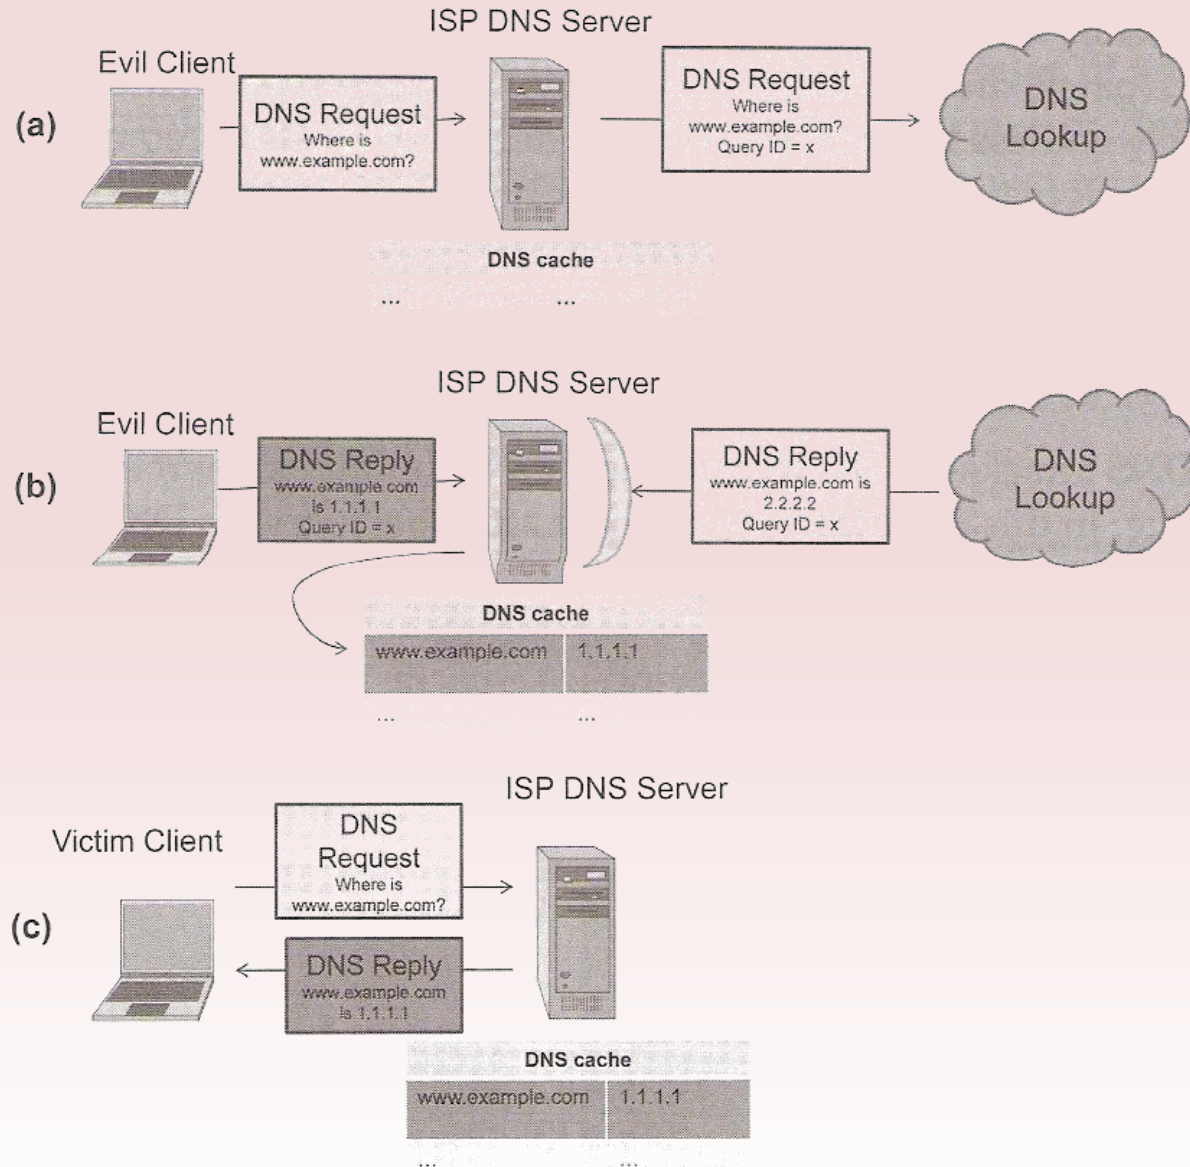
DNS Cache Poisoning. Some attacks are made possible by a technique known as DNS cache poisoning. In this technique, an attacker attempts to trick a DNS server into caching a false DNS record to that server to resolve domains to attacker-supplied IP addresses.

### DNS Cache Poisoning Scenario

1. An attacker has decided to launch a DNS cache poisoning attack against an ISP DNS server. The attacker rapidly transmits DNS queries to this server, which in turn queries an authoritative name server on behalf of the attacker.
2. The attacker simultaneously sends a DNS response to his own query, spoofing the source IP address as originating at the authoritative name server, with the destination IP set to the ISP DNS server target.
3. The ISP server accepts the attacker forged response and caches a DNS entry associating the domain the attacker requested with the malicious IP address the attacker provided in the forged response.

Now, any downstream users of that ISP will be directed to the attacker's malicious web site when they issue DNS requests to resolve the domain.

# DNS Cache Poisoning Attack



DNS Cache Poisoning and the Birthday Paradox. Unfortunately, randomization of transaction IDs does not completely solve the problem of DNS cache poisoning. If an attacker can successfully guess the ID associated with an outbound DNS request and issue a response with the same ID, an attacker could still accomplish the DNS cache poisoning attack.

This increase in attack success probability from an increase in fake requests is a result of a principle known as the birthday paradox, which states that the probability of two or more people in a group of 23 sharing the same birthday is greater than 50%.

Yielding the formula:  $23 * 22 / 2 = 253$  pairs. It only takes one matching pair of the birthday paradox to hold.

An attacker issuing a fake response will guess a transaction ID equal to one of  $n$  fake responses different 16-bit real IDs with probability  $n/2^{16}$ ; hence, it would fail to match one with probability  $1 - n/2^{16}$ . Thus, an attacker issuing  $n$  fake responses will fail to guess a transaction ID equal to one of  $n$  different 16-bit real IDs with probability

$$(1 - (n / 2^{16}))^n$$

So issuing at least  $n = 213$  requests and an equal number of random fake responses, an attacker will have roughly a 50% chance that one of the random responses will match a real request.

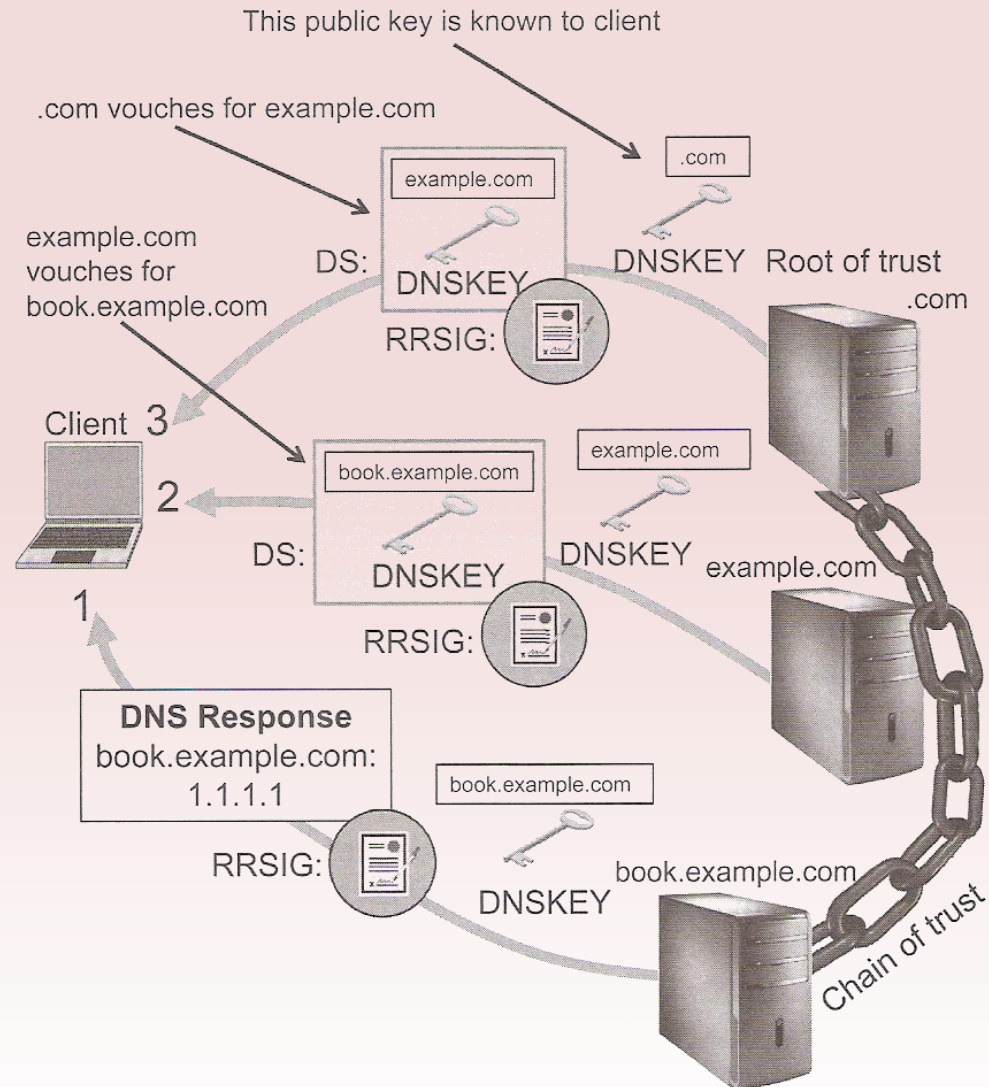
Client-Side DNS Cache Poisoning Attacks. A DNS cache poisoning attack occurs when 1) someone visits a malicious web site and views a page containing many images, each causing a separate DNS request to be made to a nonexistence subdomain of the domain that is to be poisoned. 2) The malicious web server sends guessed responses to each of the requests, and on a successful guess, the client's DNS cache is poisoned.

## DNSSEC

DNSSEC is a set of security extensions to the DNS protocol that prevents attacks such as cache poisoning by digitally signing all DNS replies using public-key cryptography. These signatures make it infeasible for an attacker to spoof a DNS reply and thereby poison a DNS cache.

To perform signature verification, the client uses the parent name server's DNSKEY to decrypt the RRSIG (Resource Record Signature) record, compares this to the DS (Designated Signer) record of the child name server's DNSKEY. This process is repeated until a "trusted key" that the client has existing knowledge of and does not need to verify is encountered. DNSSEC clients must be configured with other known trust points at levels below the root name server.

# DNSSEC Response and the Chain of Trust




# Firewalls

The internet is a vast network of untrusted and potentially malicious machines. In order to protect private networks and individual machines from dangers of the internet, a firewall can be employed to filter incoming and outgoing traffic based on a predefined set of rules that are called firewall policies.

Firewalls may be used both as a protective measure, to shield internal network users from malicious attackers on the internet, or as a means of censorship.

Ex: Many companies prevent internal users from using certain protocols or visiting certain web sites by employing firewall technology. On a larger scale, some countries, such as China, impose censorship of their citizens by subjecting them to restrictive national firewall policies that prevent users from visiting certain type of web sites.

Firewalls can be implemented in either hardware or software, and are typically deployed at the perimeter of an internal network (*at the point where the network connects to the internet*). The internet is considered an untrusted zone, the internal network is considered a zone of higher trust, and any machine(s), like the firewall, situated between the internet and the internal trusted network are in what is known as a demilitarized zone (DMZ). Firewall can also be implemented in software on personal computers.





## Firewalls Policies

Packets flowing through a firewall can have one of three outcomes:

1. Accepted: permitted through the firewall
2. Dropped: not allowed through with no indication of failure
3. Rejected: not allowed through, accompanied by an attempt to inform the source that the packet was rejected

Policies used by the firewall to handle packets are based on several properties of the packets being inspected, including:

- The protocol used (TCP or UDP)
- The source and destination IP
- The source and destination ports
- The application payload of the packet (whether it contains a virus)

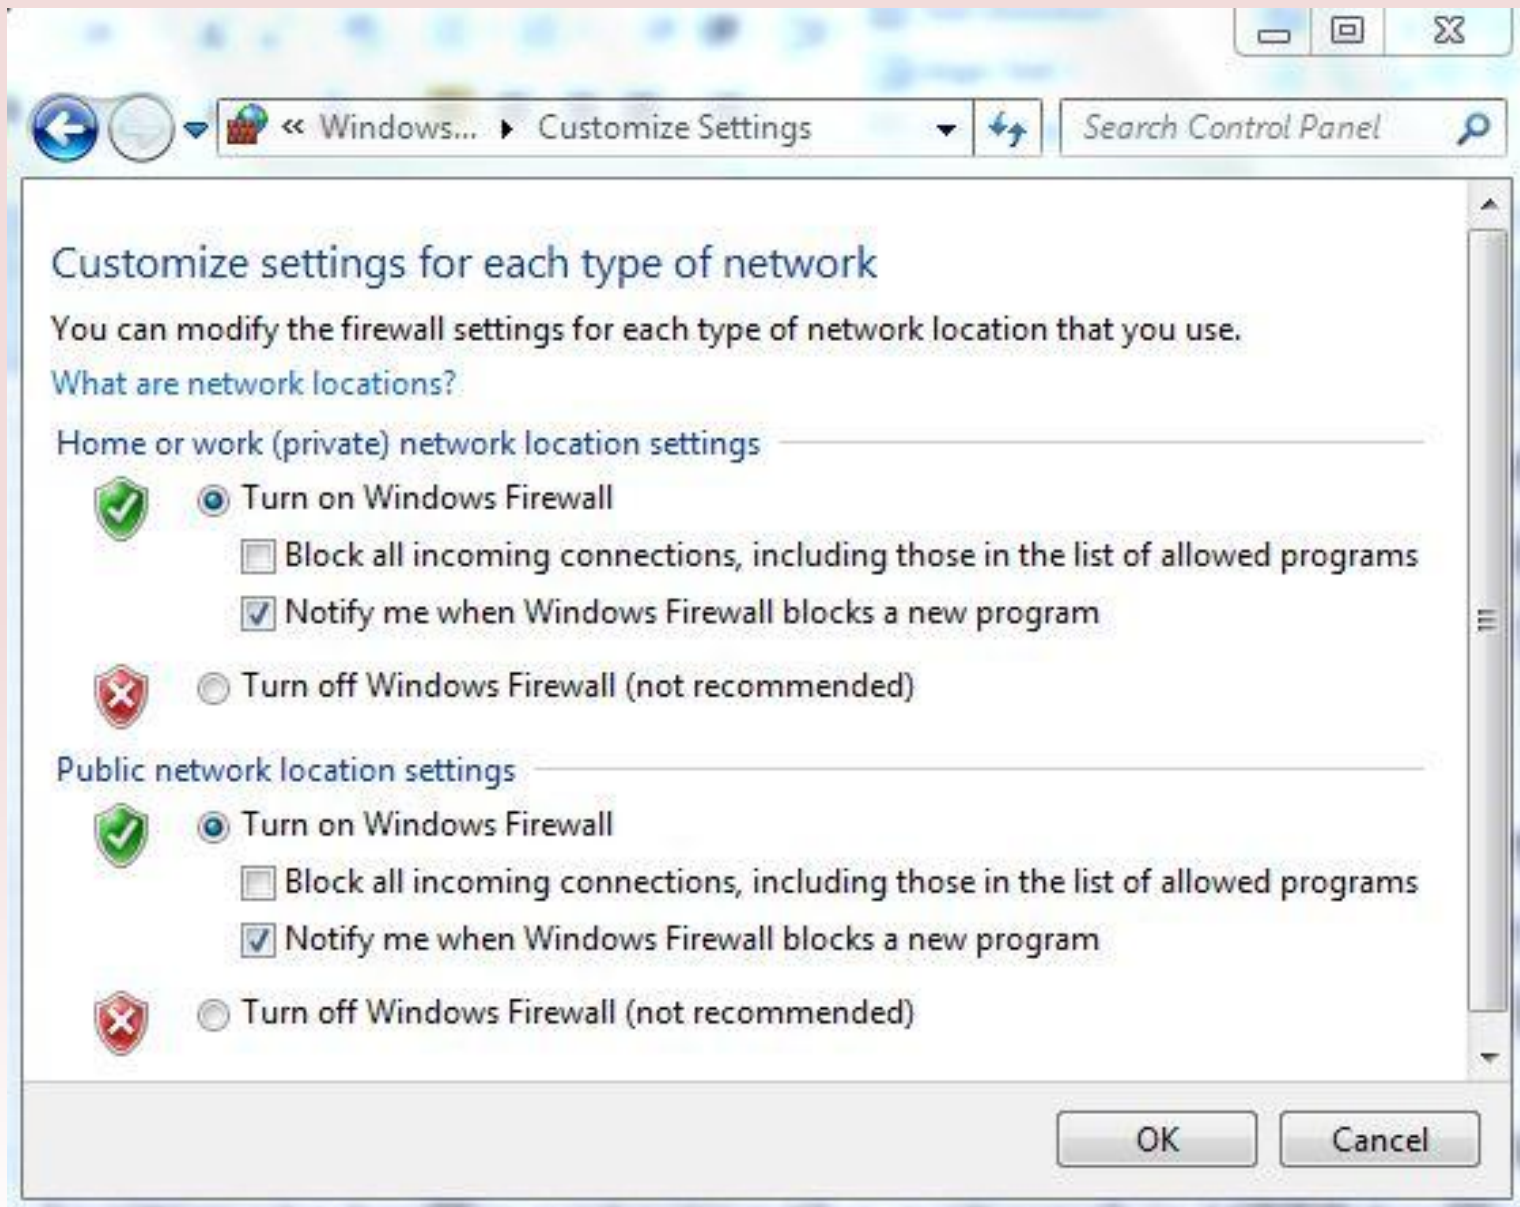
Blacklist and Whitelist. Some network administrators choose a blacklist approach, or default-allowed ruleset. In this configuration, all packets are allowed through except those that fit the rules defined specifically in the blacklist.

A safer approach to defining a firewall ruleset is to implement a whitelist or default-deny policy, in which packets are dropped or rejected unless they are specifically allowed by the firewall.

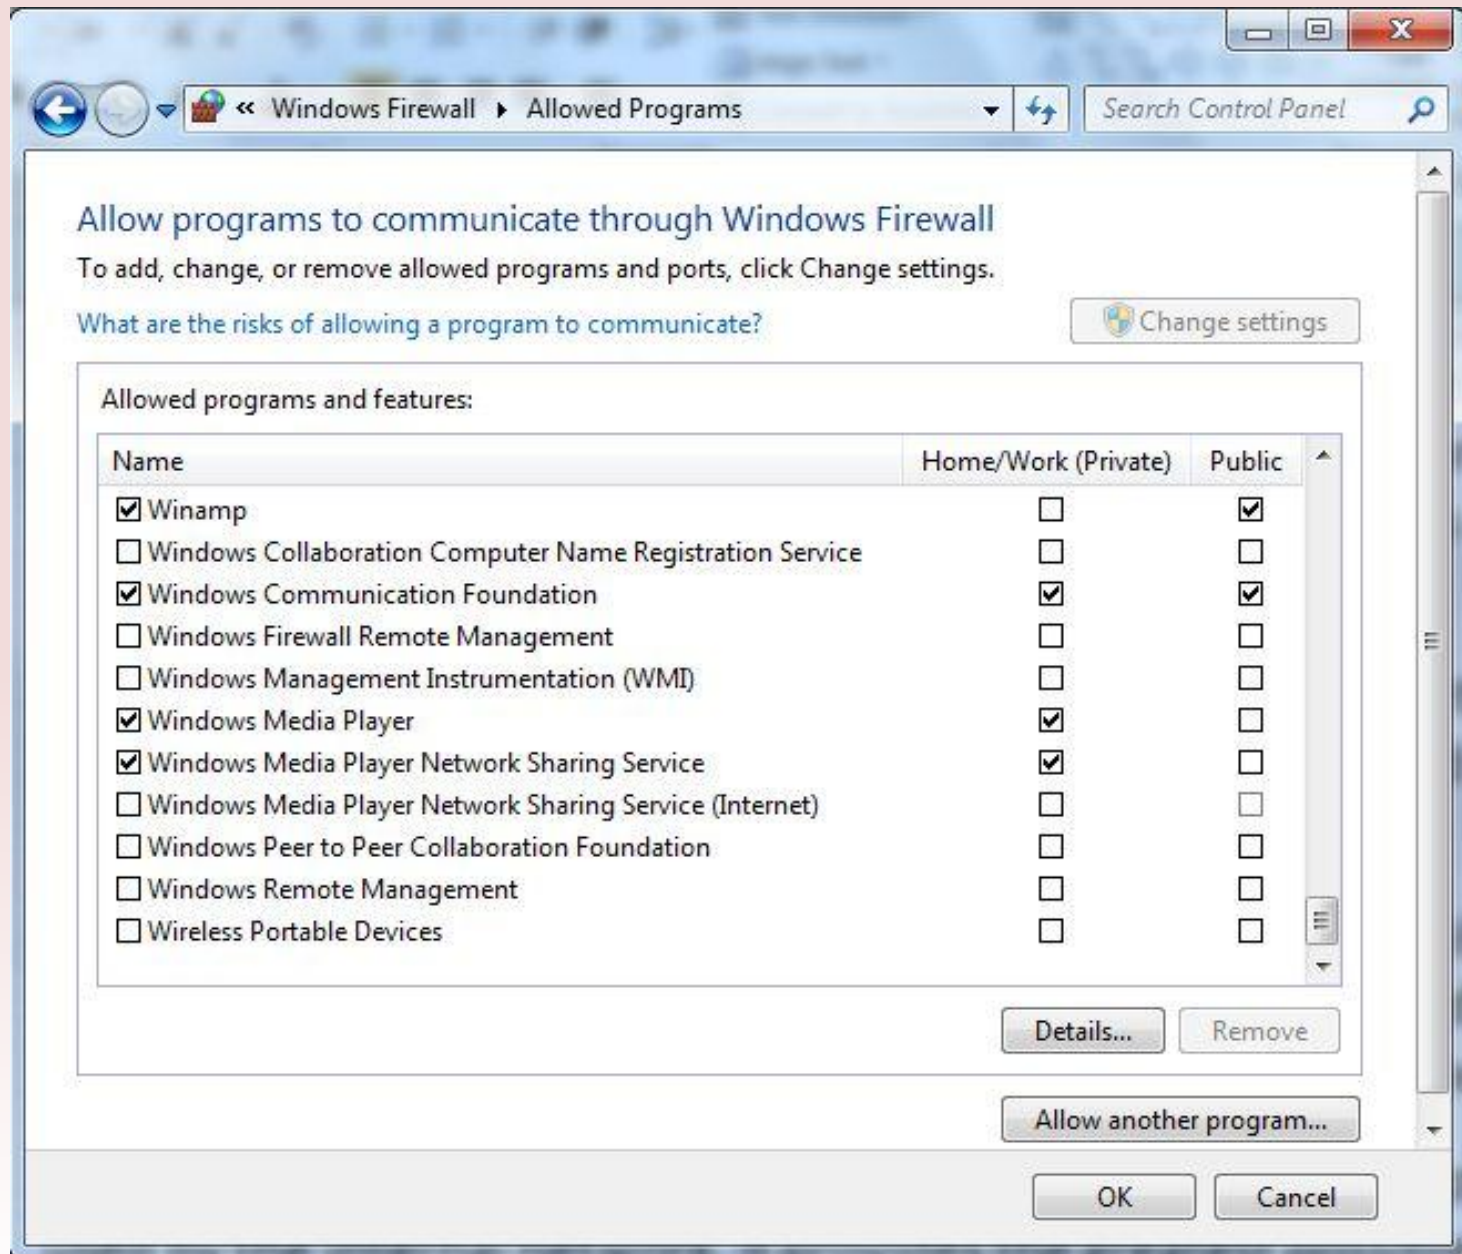
Ex: A network administrator might decide that the only legitimate traffic entering the network is HTTP traffic destined for the web server and that all other inbound traffic should be dropped.

While this configuration requires greater familiarity with the protocols used by the internal network, it provides the greatest possible caution in deciding which traffic is acceptable.

# Windows 7 - Firewall Setting




# Windows 7 - Setting Allowed Programs



# Stateless Firewalls

This implementation of a firewall does not maintain any remembered context (or state) with respect to the packets it is processing. Instead, it treats each packet attempting to travel through it in isolation without considering packets that it has processed previously.

While stateless firewalls provide a starting point for managing traffic flow between two untrusted zones and require little overhead, they lack flexibility and often require a choice between limited functionality and lax security.



## Stateful Firewalls

Stateful firewalls can tell when packets are part of legitimate sessions originating within a trusted network. Stateful firewalls maintain tables containing information on each active connection, including the IP addresses, ports, and sequence numbers of packets. Using these tables, stateful firewalls can solve the problem of only allowing inbound TCP packets that are in response to a connection initiated from within the internal network. Once the initial handshake is complete and allowed through the firewall, all subsequent communication via that connection will be allowed, until the connection is finally terminated.

Handling TCP connections is relatively straightforward because both parties must perform an initial handshake to set up the connection. Handling UDP traffic is not so clear. Most stateful firewalls consider a UDP “session” to be started when a legitimate UDP packet is allowed through the firewall. At this point, all subsequent UDP transmissions between the same two IPs and ports are allowed, until a specified time is reached.

# Stateful Firewall

## 6.2. Firewalls

