

Network Security

ICMP, TCP, DNS, Scanning

Johannes Schmidt

Institutionen för Datavetenskap (IDA)

Avdelningen för Databas- och Informationsteknik (ADIT)

Agenda


- A couple of examples of network protocols that did not have security in mind during the design.
 - **Lesson learned:** When designing new protocols it is hard to imagine all the ways the design may be exploited.
 - **Lesson learned:** It is important to learn from the past
 - **Lesson learned:** Not all development is done with security in mind, *do not blindly trust the work of others.*
- **Scanning**
 - A tool for network managers to check their configurations.
 - A tool for attackers to find potential targets.
 - Requires detailed knowledge about protocols to build, simple to use.

ICMP

- ICMP is used to send error and control messages in a network.
 - A requested service is not available.
 - Time to live exceeded in transit.
 - Echo reply (used by ping).
 - etc.
- Absolutely necessary for networks to work properly.
- Not designed for security, and has had a reputation of being a dangerous protocol.
- **Source quench:**
 - Tells the recipient to slow down sending, if obeyed, then source quench messages could be used to complete a low-bandwidth DoS.
- **Redirect:**
 - Tells a system to send traffic destined for a particular system to a specified router, bad idea as it allows DoS, MITM, etc.
- There are others, but systems today ignore the inappropriate messages.


TCP (Simplified)

- The client sends a connection request to the server.
- The server, if willing to accept the connection, responds with an acknowledgment.
- The client then acknowledges that, and the connection is established.
- **Problem:** The server upon receiving the second packet, needs to confirm that it previously accepted this connection, and that it isn't a stray packet or something bad.
- Therefore, TCP has a connection queue where all accepted connections are parked until the second packet is received.



SYN flood

- The SYN flood attack exploits this.
- A large number of connections are requested, but the second packet is never sent.
- The server allocates resources to remember the accepted connection in the queue (until time-out or second packet).
- In theory this will eat up the servers resources, making it unable to accept new requests.
- In practice this is not much of a problem anymore, ***but when developing new protocols it is easy to miss these details.***



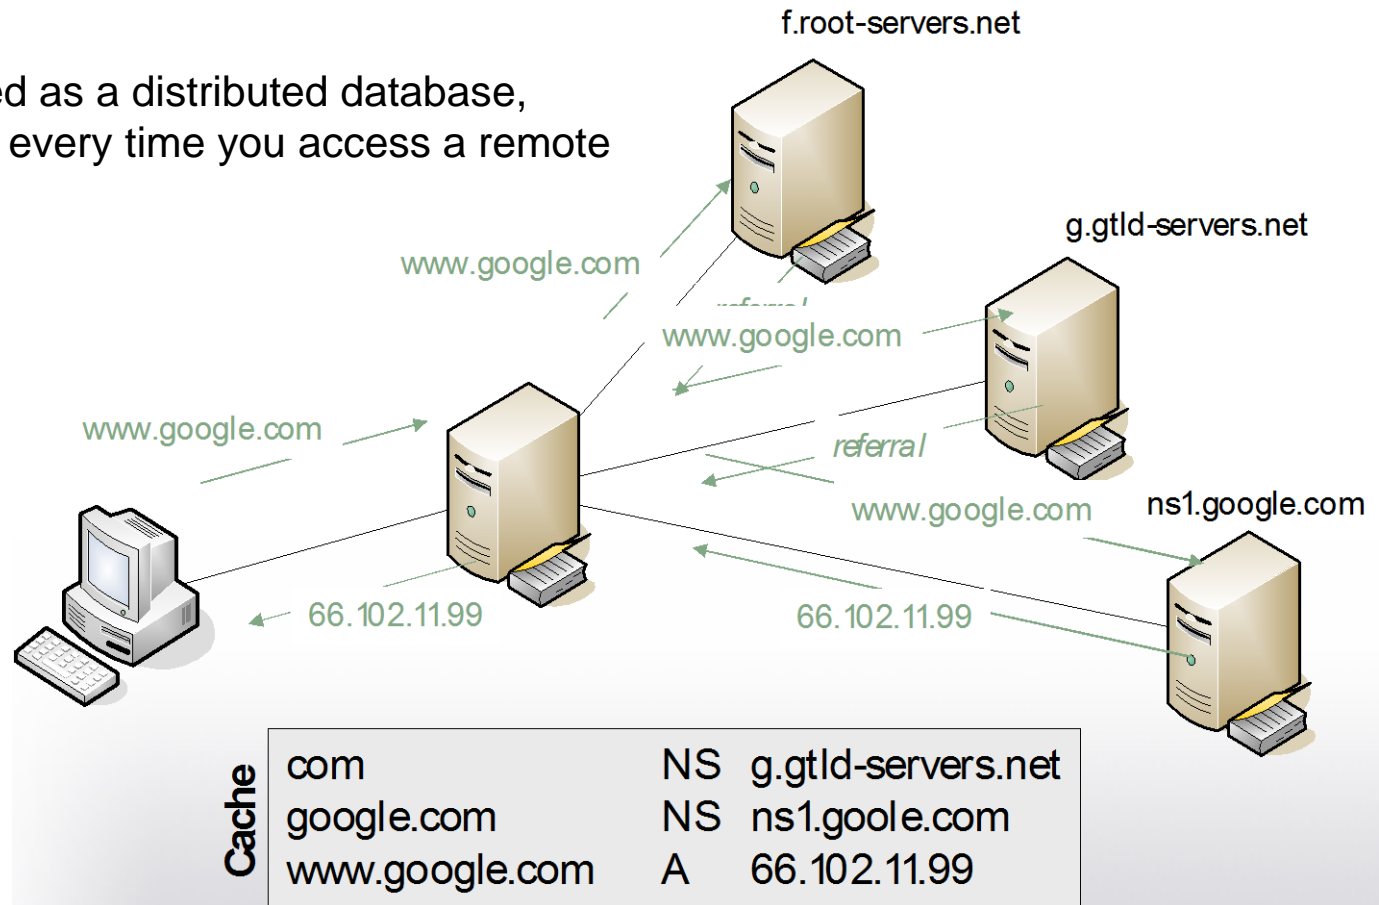
Preventing SYN floods

- One way of preventing SYN flood attacks is to not save any information at all on the server about accepted connections.
- Instead a *cookie* is sent along with the acknowledgement from the server. This cookie contains information only the server knows.
- The client attaches the cookie to its response, and the server is able to verify that the cookie is correct, and that the client has previously been accepted.
- An attempt has been made to add this to TCP, but it turned out to be quite ugly. But for new protocols it has been implemented successfully (SCTP).
- ***Learn from the past to create better solutions in the future.***

DNS Security

An absolute vital service for the Internet as we know it. Primary purpose to map easy-to-remember names to addresses.

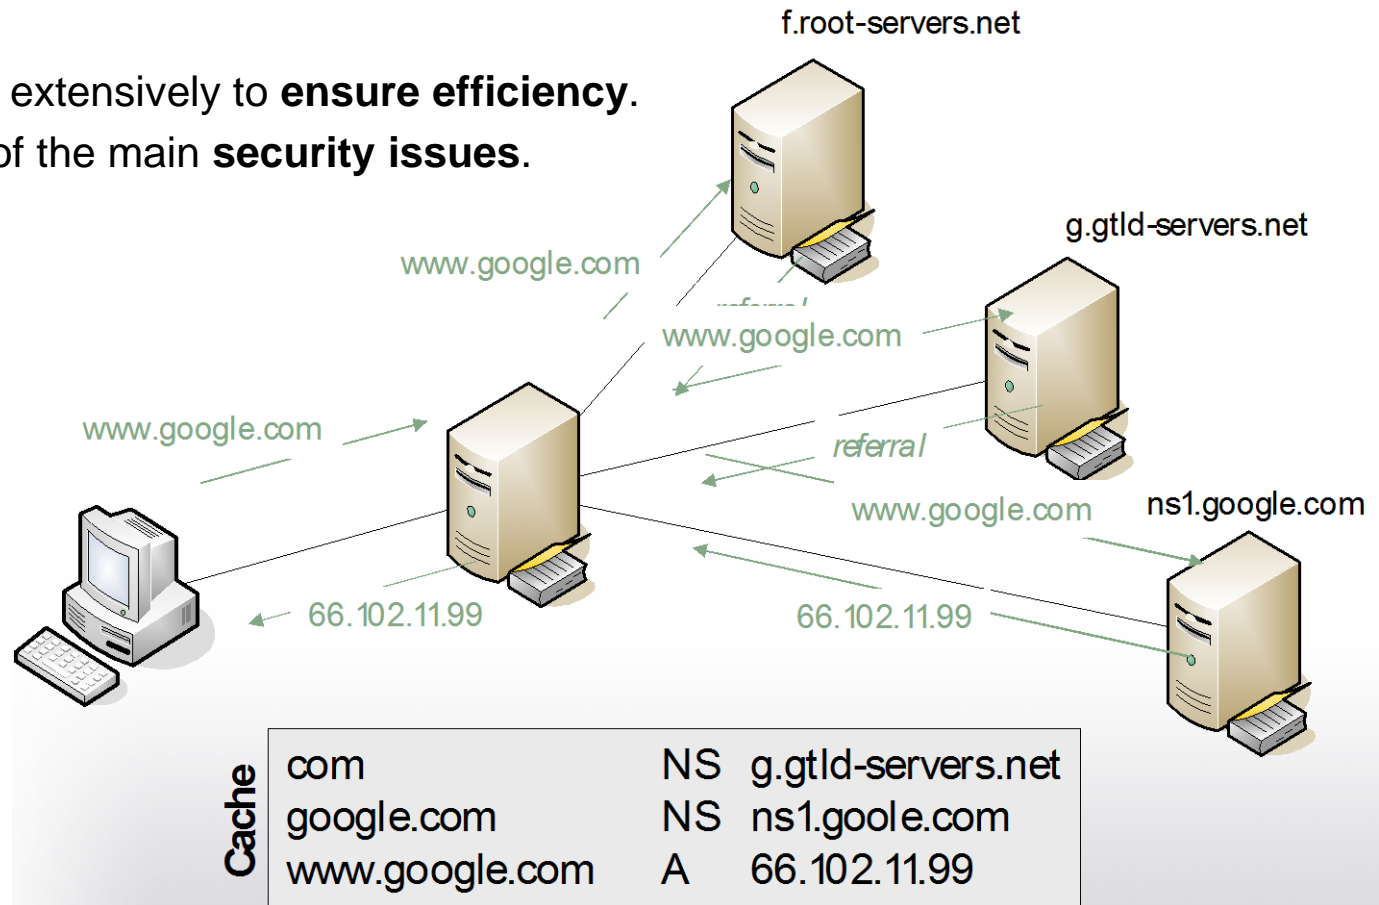
DNS can be viewed as a distributed database, which is accessed every time you access a remote system.



DNS Security

The full database is distributed over a huge number of hosts, ordered in a hierarchical fashion, facilitated to find information rapidly.

- **Caching** is used extensively to **ensure efficiency**.
- **Caching** is one of the main **security issues**.

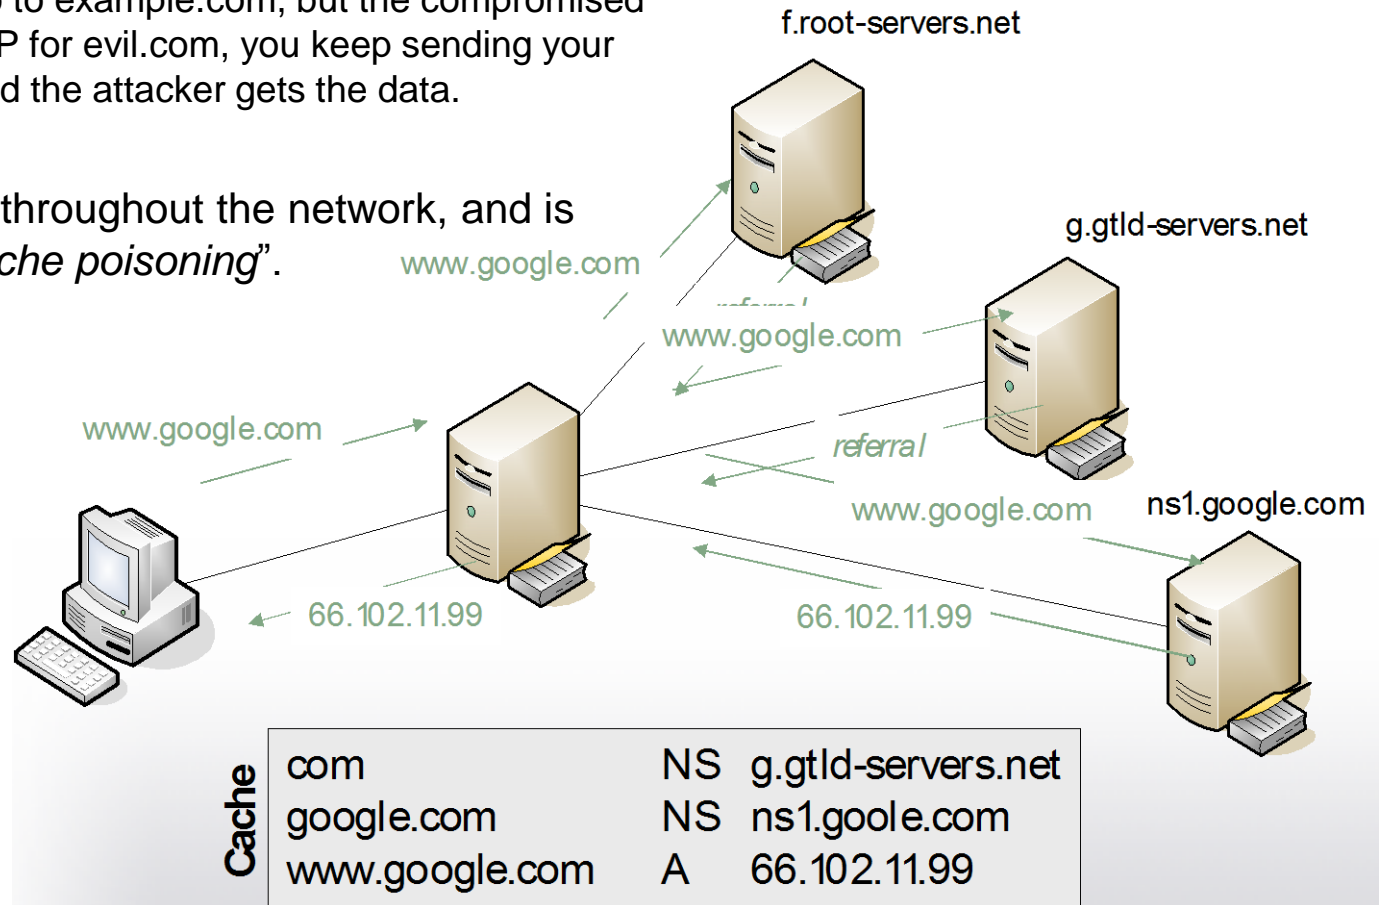


DNS Security

What if an attacker compromised the DNS database?

Example: You want to go to example.com, but the compromised DNS responds with the IP for evil.com, you keep sending your private data to this IP, and the attacker gets the data.

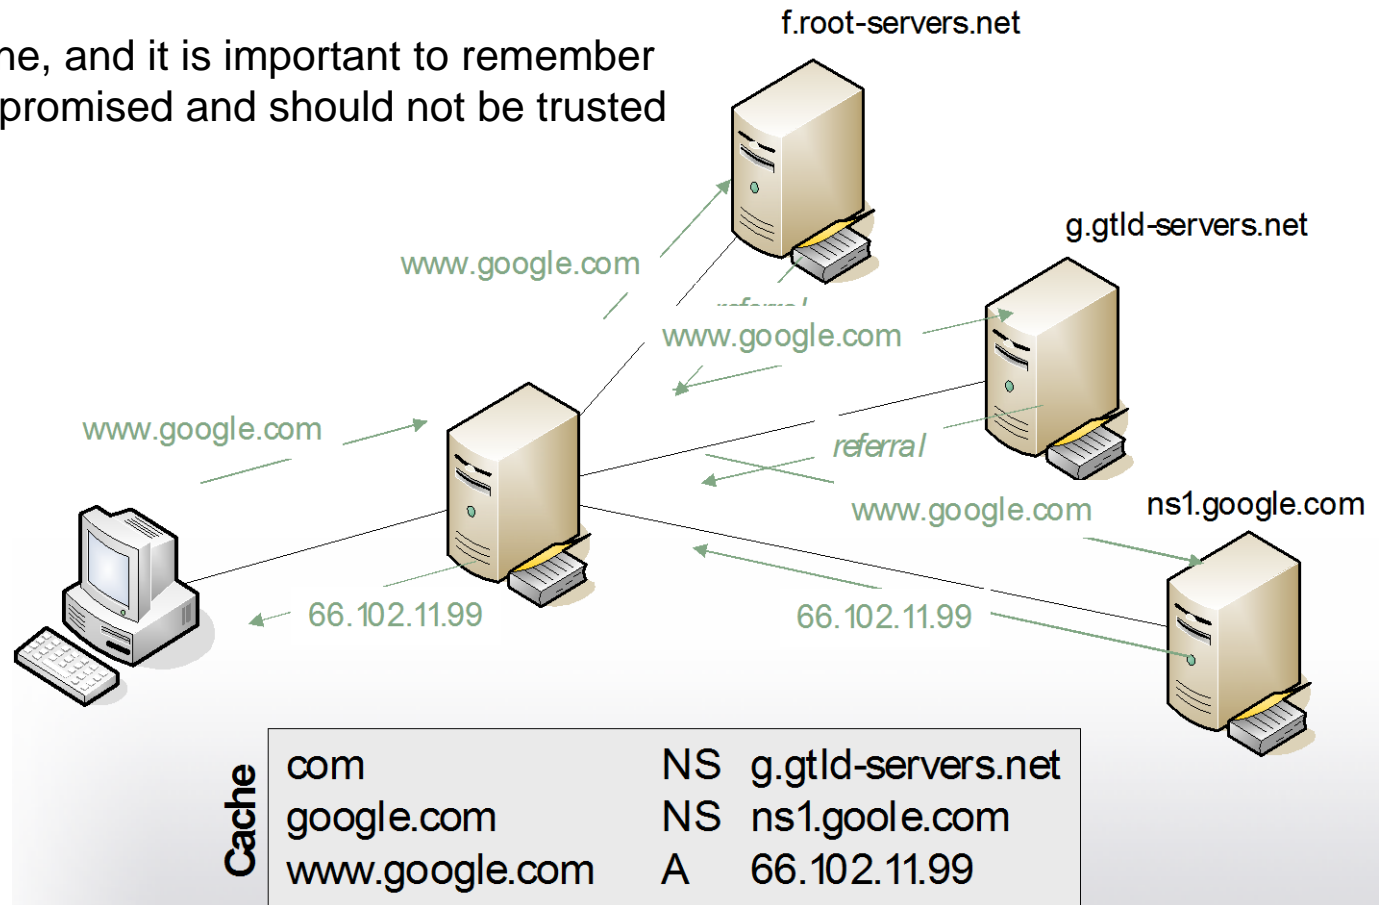
This also propagates throughout the network, and is cached, so called “*cache poisoning*”.



DNS Security

Cache poisoning has become harder since new implementations protect against it.

However it can be done, and it is important to remember that DNS can be compromised and should not be trusted blindly.



SCANNING

What do people scan for?

- Firewall configurations
 - Open ports
 - RPC services
 - Known vulnerabilities
 - Operating system details
-
- The purpose of the scan is to gather information that can be used later in an attack.
 - There are tools that greatly simplify the job, such as *nmap*.
 - *nmap* will determine open ports and operating system details.
 - *Firewalk* is another tool for firewall configuration scanning.

What do people scan for?

- Building a scanner requires very good understanding of all the parts in a system.
- One has to know the differences in response if a request is sent to an open port or not, etc.
 - These differences often depend on the used implementation and the configuration of the system.
 - There are differences depending on the operating system version, firewall configuration, etc.

Firewalking

The goal is to figure out the configuration of a firewall without connecting to any system beyond it.

Theory:

- Most firewalls will inspect each datagram before sending it to the forwarding engine of the firewall.
- If a datagram is not passed by the firewall rules, it is usually discarded (other reactions can be used).
- If it passes the rules, it is processed for forwarding, which involves decrementing its TTL.
- If the TTL reaches zero then an *ICMP destination unreachable/time exceeded* is sent.

Firewalking


The goal is to figure out the configuration of a firewall without connecting to any system beyond it.

Practice:

- Firewalking works by sending datagrams that have a TTL set such that if it is passed by the firewall it will be zero.
- The datagrams are sent to any address behind the firewall.
- If the firewall does not respond, this means the port is filtered; if an *ICMP destination unreachable/time exceeded* is received then the port is open.


TCP Scans

- Knowing a port is open in a firewall is useful, but knowing if there is a process actually listening to this port is even *more* so.
- TCP SYN scan**
 - Send SYN segments to the server. If a SYN/ACK is sent back then the port is used, immediately send a RST back (this tears down the connection before it is complete, and usually avoids logging).
 - If server responds with RST then no process on port.
 - If no response, then firewall is probably filtering.



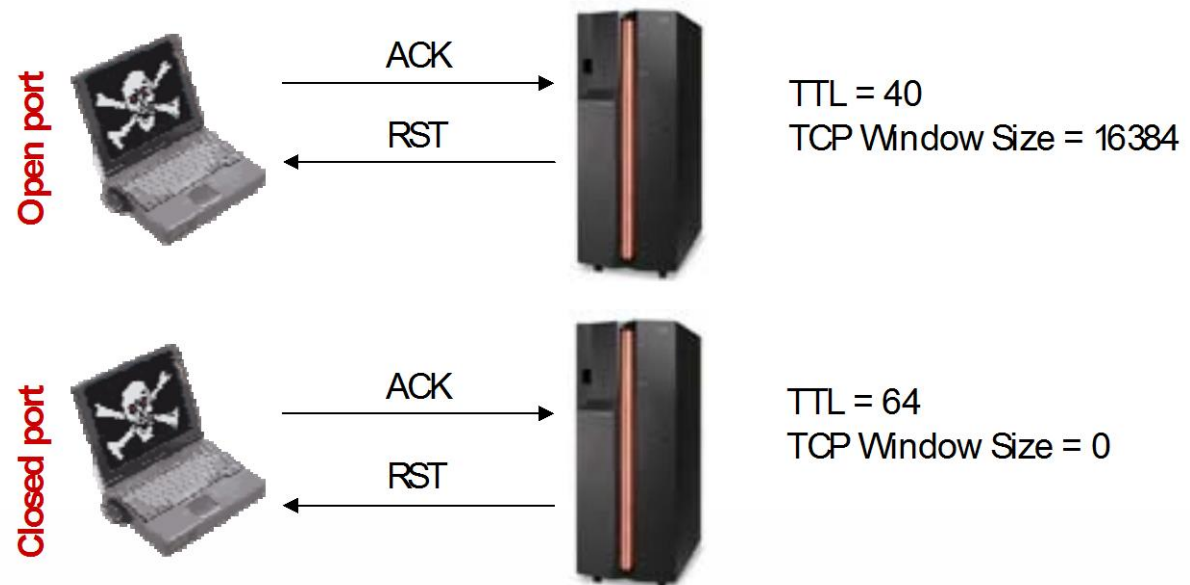
TCP Scans

- Alternative to SYN scan is FIN scan.
 - Attacker sends segment with FIN flag.
 - Firewalls that only filter *connection* attempts will let all FIN segments through.
 - TCP specification requires that an RST is responded if destination port is closed, and ignores the FIN if it is open.




TCP Scans

- Implementations of TCP can behave differently.
- A third scanning version called ACK scan
 - RST is returned regardless if the port is used or not.
 - But there are then instead differences in the TTL and window size in the response, these can then be read and used instead.



UDP Scans

- UDP scanning is a bit harder than TCP scanning.
- Depends on the target responding with a “port unreachable” message when a UDP datagram is sent to a closed port.
- When sent to an open port the response is unpredictable, sometimes nothing is done at all. This makes it hard to distinguish between a filtered and open port.



Summary

- **Lesson learned:** When designing new protocols it is hard to imagine all the ways the design may be exploited.
- **Lesson learned:** It is important to learn from the past
- **Lesson learned:** Not all development is done with security in mind, do not blindly trust the work of others.
- **Scanning**
 - A tool for network managers to check their configurations.
 - A tool for attackers to find potential targets.
 - Requires detailed knowledge about protocols.



Linköpings universitet

expanding reality

www.liu.se