

CSN09101

Networked Services

Week 12: Hacking

Module Leader: Dr Gordon Russell

Lecturers: G. Russell

This lecture

- Hacking Techniques
- DNS cache poisoning
- Discussions

Hacking Techniques

Security Concerns

- Security matters to most people
- You want to stop people doing things you do not want them to do with your resources.
- There are different approaches to security.
 - Secure things when problems occur
 - Proactive protection measures

Big Danger

- There is a real danger in providing secure environments – they will become less usable to the real users...
 - Example – stop viruses spreading by blocking outgoing SMTP stops people sending emails without a proxy.
- If users think your security is in the way, they will take steps to bypass your security.
- The ideal security is invisible to real users, but unavoidable to hackers.

Hacking

- Perhaps the best way to understand security is to understand hacking.
- There are black hats, white hats, and some in between...
 - Black hats break systems maliciously
 - White hats break systems without damaging things, and help admins become more secure.
 - Grey hats tend to break “some” systems maliciously, like pornography sites, or break things accidentally, like students doing coursework...

Cost

- All proactive measures, and all active hacking attempts incur an administrative cost.
- Hat colour does not reduce cost.
- Proactive measures are hard to cost-justify.
 - “Yes boss, I did spend a week securing our systems. You wont notice the change, but we are 30% more secure now...”
- Reactive measures are hard to tolerate.
 - “Yes maybe I should have taken security seriously, but how was I to know we were going to be hacked? We have never been hacked before!”

The approach

- There are three approaches...
 1. Social engineering
 2. Brute force
 3. Technical intrusions
- Most attacks are about escalation of privileges.
 - Do not necessary attack ROOT, but get some sort of privilege on the remote machine and use that to become more powerful.
 - Hacking into a basic account then getting more privileges is know as privilege escalation.

Social Engineering

- This is becoming very common.
- Relies on human nature.
- Good examples are:
 - Emails from your bank asking for username and password details.
 - Phone calls from an “administrator”.
 - Visits from offsite technicians.

Hacking BLOB University

- Firstly I get the name of a user and their telephone number.
- I also need a pay as you go phone.
- I pretend to be someone in support.

\$ whois blob.ac.uk

...

Registrant Contact:

Jimmy Smith

Registrant Address:

Director of IT Services

Blob University

Blobby House

+44 141 555 5555

Hacking BLOB University

- Its best if Jimmy is not in when you phone. But a Director will never answer their phone anyway...
- Phone your target, “Hi, Jimmy Smith here from Blobby House, your user account may have been hacked as your account has sent obscene messages to a secretary in accounts. Unless you want to admit to it now? Am I talking to the right person? Whats your userid? Blah blah. For auditing reasons, you better stop using your account for the rest of the day. What’s the password, as I need to supply that to the police.”

As strong as the weakest link

- Visit some universities at the start of the year, and you may even see their password policy.
 - “Your surname plus the year of your birth”.
- Passwords for chocolate:
 - <http://news.bbc.co.uk/1/hi/technology/3639679.stm>
 - When surveyed on the street by attractive “surveyors”, 70% of people gave up their password for a bar of chocolate. In fact, 34% of those questioned gave their password without the offer of chocolate...

Brute Force

- There are plenty of software tools to do this.
- An example of this could be a password cracker testing passwords from a dictionary to gain user access to a site.
- Another example is Denial of Service, with the sheer number of something causing performance degradation.

DoS

- Denial of service is popular.
- Stopping a site doing its work costs real and calculatable amounts of money.
- If a gaming site takes in 100,000 per day, then it may be justifiable to pay someone 50,000 to stop them bringing it down for days in a row.
- Simple DoS is easy to protect against, as all you need is more bandwidth than the attacker.

Example: SMURF

- If you sent a PING to a machine, it pings back a response.
- If you PING a machine, but forge the SRC address to someone you don't like, this other machine gets the PING replies...
- On some systems, a PING to a.b.c.0 (the network) will get all machines on the network to PING back.
- Combine this together and you have an untraceable DoS attack!

Distributed DoS

- If a hacker has taken over a few thousand computers using a trojan or virus, they can build themselves a bot farm.
- On command each machine can launch a hard-to-trace DoS attack on you.
- This can be much harder to block
 - You can block the src address but if this is forged you are stuffed.
 - Best you can do is rate limit syn packets, so that at least legitimate connections are treated normally after the tcp handshake.

Technical Exploits

- Exploit deficiencies in system design, configuration, or management.
- Most involve 5 problem areas:
 1. Inherent security defects
 2. Misuse of legitimate tools
 3. Improper maintenance
 4. Ineffective Security
 5. Inadequate detection systems

Security Defects

- Software is now so complex that all software ships with with unexpected “features”.
- Problems are often reported publicly (e.g. on CERT, SANS, CVE).
- Vendors will eventually release “fixes”.
- Some time later system admins will install the new version...
- Not all problems can be fixed (e.g. protocol weaknesses).

Misusing Tools

- Many useful tools in standard installs can be used to break security if misused..
 - ping – find victims
 - traceroute – find network topologoes.
 - dig – DNS information
 - whois – background information on target.
 - finger – who is logged in.
 - rpcinfo – what rpc services are running
 - showmount – what NFS mounts are exported
 - telnet – play with any TCP protocol service.

> showmount -a orion.napier.ac.uk | grep gor

```
artemis:/export/home/o2/staff/gor
```

```
pc236nt:/export/home/o2/staff/gor
```

> mount -t nfs orion.napier.ac.uk:/export/home/o2/staff/gor
/mnt/a

- I have access to all I: drive files for all users.
- No password required.
- Moral: do not use I: for your vital stuff...

Improper Maintenance

- An example of this could be firmware in a router not be updated, or critical updates to a system being missed.
- Lack of time is often the cause.
- Lack of priority may also be an issue.
- Perhaps highlights the need to prioritise security matters as determined by a risk assessment.

Ineffective Security

- Perhaps caused by having no security policy...
- An example could be spending all of ones time on firewall management, and then leaving the root password as “rootroot” for speed.
- Also causes by conflicts between users and admin people.
 - Users want CGI and write scripts which bypass admin security.

Detection

- Many sites rely on audit trails to detect problems.
- This does nothing to detect Trojans, backdoors, and viruses.
- New tools on the market to detect more subtle problems.
 - For example, checksums of system files checked against remote records.

The Process

- A dedicated hacker will have many targets on the go at once, all in different stages of being hacked.
- A commonly held list of stages is:
 1. Casing
 2. Scanning
 3. Enumeration

Casing:

- Gather information on the target.
- Often called fingerprinting.
 - IPs, services running, routing tables, domain information, authentication scheme, user details, admin names, contact information, telephone numbers, connection type, etc.
 - For instance, has an admin discussed their firewall configurations on a newsgroup or admin forum? Could be interesting...

Scanning

- In the scanning phase, individual machines are identified by direct communication.
- Tools are available to tell you about OS type, open ports, firewall configurations, and even version numbers.
- Scanning should also involve the routers and firewall devices, as these may be remotely configurable.

Example: nmap

> nmap linuxzoo.net

PORT	STATE	SERVICE
22/tcp	open	ssh
23/tcp	open	telnet
53/tcp	open	domain
80/tcp	open	http
81/tcp	open	host2-ns
123/tcp	closed	ntp
5900/tcp	closed	vnc
5901/tcp	closed	vnc-1
5902/tcp	closed	vnc-2
5903/tcp	closed	vnc-3

Enumeration

- This really covers getting some sort of “access”.
- It could be discovering a username and then a password (brute force perhaps).
- It could be a badly written NFS or other share.
 - NULL Shares
 - Zone transfers

Failed SSH logins

- /var/log/secure contains SSH attempts.
- Failed attempts can be a sign of a problem:

```
head -3 /var/log/secure
```

```
Failed password for root from ::ffff:219.232.?.? port 40731 ssh2
```

```
Illegal user eaguilar from ::ffff:61.129.?.?
```

```
Failed password for illegal user eaguilar from ::ffff:61.129.?.? port 53785 ssh2
```

- Not exactly readable... “?” inserted to spare embarrassments.

A little perl

```
#!/usr/bin/perl

open(my $file,"</var/log/secure.1");
my %ip;
foreach my $line (<$file>) {
    if ($line =~ m/Failed/) {
        if ($line =~ m/(\d+\.\d+\.\d+\.\d+)/) {
            $ip{$1}++;
        }
    }
}

foreach my $k (keys %ip) {
    my $what = `dig -x $k | grep PTR`;
    my $dig = "?";
    $dig = $1 if ($what =~ m/PTR\s+([\s]+)$/);
    print "$k \t: $ip{$k} fails : $dig\n";
}
```

The data

202.118.?.? : 73 fails : ?
140.125.?.? : 1 fails : ?.?.edu.tw.
61.129.?.? : 23 fails : ?
78.110.?.? : 239 fails : ?
152.104.?.? : 24 fails : static-ip-?.rev.dyxnet.com.
219.232.?.? : 1 fails : ?

- Not everyone has a PTR record. The 78.110.?.? has a lot of attempts.
- Even without PTR, we can use WHOIS or RIPE to find out.. It comes to a dedicated server company in the UK.

Next Step

- Cannot really take action.
- Even UK site is probably the result of being hacked itself.
- Some people have these scripts add iptables rules so that more than “n” failed logins blocks the IP for a certain period of time.

Any successful login is the end...

- Once a hacker can actually run commands on a machine, the end is likely to be close...
- There are a great many bugs in normal everyday system commands which allows people to become root, access things they are not allowed to, and to modify things we would rather they didn't.
- Security ends with a successful login.

DNS Cache Poisoning

DNS Cache Poisoning

- In 2008 there were plenty of security issues.
- One of the big ones was to do with DNS Cache Poisoning. Also known as the Kaminsky DNS Vulnerability.
- A easy-to-read summary is at:
<http://www.unixwiz.net/techtips/iguide-kaminsky-dns-vuln.html>
- As you know, your DNS queries are probably cached in your ISP Nameserver.
- DNS Queries themselves are usually just a single packet, and the reply is probably a single UDP packet too

Messing up the cache

- If you knew that the nameserver had just sent a query for “linuxzoo.net”, you could send a forged reply BEFORE the real reply was sent.
- In the forgery you could put in your own NS, A records, or glue.
- If your reply was received before the real one, then the real reply would be ignored.
- So a hacker asks a nameserver to look up linuxzoo.net, then fires in a forged reply.

Transaction ID

- The people who came up with DNS must have thought there would be a problem, so put in a Transaction ID in each query.
- The ID of the query must match the ID in the reply in order for it to be accepted.
- So to fake a reply you have to produce the right ID.

Faking the ID

- There are 2 approaches to this.
 - On some nameservers the ID is simply incremented by one each time. So a hacker get the system to lookup something involving an authoritative nameserver they control, capture the packet and its ID. Then they ask for “linuxzoo.net” and forge a reply with the ID+1 (and +2, +3, and a few more just in case they missed a query). Easy.
 - Better systems use a random ID. To forge that you need to try and guess the ID.

Guessing the ID

- If you are lucky you can send 50 forged guesses before the real answer is received.
- The ID is only 16 bits in size. 1 in $2^{16}/50$ of being right. That's actually not bad odds.
- Previously the thought was that you only get one shot at this hack, as after that the right answer is in the cache so there is no more queries and thus no reply to forge...

The twist

- However, the twist is that each time you try to forge a packet, you ask the ISP nameserver for a random host in each domain, e.g.
 - Blahblah1234.linuxzoo.net
 - Blahblah1235.linuxzoo.net
- As each query is different, it won't be in the cache. You get another try each time you make the query.
- The clever bit is that in the forged reply you don't give the A record, you give a delegate reply. This overrides the current data in the cache, and transfers the whole domain to what your delegation information was in your forged reply. You now own the domain...
- Apparently you can take over a domain in under 10s.

The Fix

- This problem is a fundamental problem in DNS.
- Fixes are hard without rewriting DNS.
- The quick fix is that the IP nameserver uses a random transaction ID AND a random source port number. You have to guess both. This guess is so hard that excessive (and easy to spot) traffic is generated for 10 or more hours in a row...
- Ideally you want a bigger transaction id. That breaks the internet without everyone updating.
- DNSSEC may fix it, by “signing” DNS requests. That’s years away.

Discussion

Discussion

- In a DNS Cache Poisoning attack, it has been suggested that the SSL certificates used by https means that banking is still secure even if the DNS cache is poisoned. Discuss.

Discussion

- Here are some mock exam questions you should now be able to answer:

Question 1

You have detected 15 failed ssh logins from 10.0.0.1. What should you do next and why?

Question 2

A firewall audit using nmap was performed, and produced the following results:

PORT	STATE	SERVICE
22/tcp	open	ssh
23/tcp	open	telnet
53/tcp	open	domain
80/tcp	open	http
123/tcp	closed	ntp

Comment on the quality of your server security.