Edinburgh Napier UNIVERSITY

CSN09101
Networked Services
Week 11: Email Management

Module Leader: Dr Gordon Russell
Lecturers: G. Russell

Edinburgh Napier UNIVERSITY

This lecture

- SMTP
- Linux Email
- Discussions

Edinburgh Napier
UNIVERSITY

SMTP

Edinburgh Napier UNIVERSITY

SMTP

- Email is send between source and destination using a simple protocol.
- SMTP is the basic protocol used
  - Simple Mail Transport Protocol (RFC 821 and 1123)
- SMTP deals mostly with simple email without attachments.
- Attachment emails (plus email delivery using some advanced features) uses ESMTP
  - Extended/Enhanced SMTP (RFC 1869).

Edinburgh Napier UNIVERSITY

SMTP

- Very simple protocol
- Text Based
- You can send email using TELNET.
- Easy to forge an email…

Edinburgh Napier UNIVERSITY

telnet grussell.org 25

Connected to grussell.org (10.10.5.5).

Escape character is '^]'.

220 grussell.org ESMTP Sendmail 8.12.11/8.12.11; Sun, 14 Nov 200819:01:01 GMT

> **helo pc236b.napier.ac.uk**

250 grussell.org Hello pc236b.napier.ac.uk [10.4.5.6], pleased to meet you

> **mail from: g.russell@napier.ac.uk**

250 2.1.0 g.russell@napier.ac.uk … Sender ok

> **rcpt to: me@grussell.org**

250 2.1.5 me@grussell.org.. Recipient ok

Edinburgh Napier UNIVERSITY

telnet grussell.org 25

> **data**
354 Enter mail, end with "." on a line by itself
> From: "Santa" claws@northpole.com
> To: Gordon Russell <me@grussell.org>
> **Date: Tue, 15 Jan 2008 16:02:43 -0500**
> Subject: SMTP

Hello gordon.
I am Santa
> .
250 Ok: queued as 5555
> QUIT
221 Bye

Edinburgh Napier UNIVERSITY

Envelope and Headers

- The email information about from and to supplied using SMTP (except the contents of the DATA command) makes up what is known as the Message Envelope.
- The from, to, and other initial information in the email itself (the part in DATA) is known as the email header.
- The email instructions in the envelope DOES NOT have to match that in the email headers.
- This is useful when, for instance, dealing with email mailing lists: The envelope directs it to you, but the headers state it is actually to “the name of the email group”.

Edinburgh Napier UNIVERSITY

Forged Emails

- Note that the envelope FROM was:
  > **mail from: g.russell@napier.ac.uk**

- The data FROM was:
  From: “Santa" claws@northpole.com

- This is perfectly valid. It will be delivered.
- The envelope is used through the delivery process, but it is discarded when it is finally delivered to the recipient.
- The final recipient cannot recover the information in the envelope. However, the headers can give useful information.

Edinburgh Napier UNIVERSITY

Email headers

From g.russell@napier.ac.uk Sun Nov 15 11:12:21 2009
Received: from pc236b.napier.ac.uk [10.2.4.5]
by grussell.org (8.18.11) id PDQ666
Sun Nov 15 11:12:20 2008 -0000
Received: (gor@localhost)
by pc236b.napier.ac.uk (8.18.11) id LXY123
Sun Nov 15 11:12:16 2008 -0000
Date: Sun Nov 15 11:12:15 2008 -0000
From: g.russell@napier.ac.uk
To: me@grussell.org
Message-Id: <20041115111215.LXY123@pc236b.napier.ac.uk>
Subject: Wow

Message body is here. This is the message.

Edinburgh Napier UNIVERSITY

- From “Date” down is the data added by the original sender.
- As the email moves from machine to machine, extra information is added to the data.

```
From g.russell@napier.ac.uk Sun Nov 15 11:12:21 2009
Received: from pc236b.napier.ac.uk [10.2.4.5]
        by grussell.org (8.18.11) id PDQ666
        Sun Nov 15 11:12:20 2008 -0000
Received: (gor@localhost)
        by pc236b.napier.ac.uk (8.18.11) id LXY123
        Sun Nov 15 11:12:16 2008 -0000
```

Edinburgh Napier UNIVERSITY

Received

Received: from pc236b.napier.ac.uk [10.2.4.5]
by grussell.org (8.18.11) id PDQ666
Sun Nov 15 11:12:20 2008 -0000
Received: (gor@localhost)
by pc236b.napier.ac.uk (8.18.11) id LXY123
Sun Nov 15 11:12:16 2008 -0000

- The first “hop” the email went through was at pc236b.
- The email was written by someone on that machine (user gor).
- Sendmail handled that hop, version 8.18.11

Edinburgh Napier UNIVERSITY

Received

Received: from pc236b.napier.ac.uk [10.2.4.5]
by grussell.org (8.18.11) id PDQ666
Sun Nov 15 11:12:20 2008 -0000
Received: (gor@localhost)
by pc236b.napier.ac.uk (8.18.11) id LXY123
Sun Nov 15 11:12:16 2008 -0000

- The second “hop” the email went through was at grussell.org. It received the email from a server with ip 10.2.4.5.
- Pc236b and 10.2.4.5 should be the same thing.
- It took 4 seconds to be delivered between servers.
- Sendmail handled that hop, version 8.18.11

Edinburgh Napier UNIVERSITY

Spotting forged emails

- You are looking for “funnies” in the headers.
  - Dates and times that go backwards (taking into account the timezone)
  - Hops which don’t match up
  - Strange strings in the hop data
  - DATA and hop data which makes no sense.
  - HOP routes which sound strange (like a bank delivering emails via yahoo).

Edinburgh Napier UNIVERSITY

MX Records

- When you email “linuxzoo.net” the delivery process will look for a MX record for linuxzoo.net.
- If it doesn’t find one, email is directed to the A record.
- If it finds a MX record, email id delivered to the machine described in the MX record.
- This allows a whole domain to delegate email reception to one or more key servers, without having to have email servers on every single possible host.

Edinburgh Napier
UNIVERSITY

Linux Email

Edinburgh Napier UNIVERSITY

MUA, MTA, and MDA

- Email in Linux is controlled via three types of services:

1. MUA – Mail User Agent
2. MTA – Mail Transfer Agent
3. MDA – Mail Delivery Agent

Edinburgh Napier UNIVERSITY

MUA

- The email “client”.
- Users use the Mail User Agent to read and send emails.
- It takes email messages which have been delivered to a particular user’s mailbox and displays them to the user.
- It takes new messages and passes these to the MTA for delivery.
- Examples include mutt, mail, and pine.

Edinburgh Napier UNIVERSITY

MTA

- The Mail Transfer Agent is the mail equivalent to an IP router.
- It takes messages given by an MUA or another MTA, and depending on the delivery address passes them onto another MTA or to an MDA for delivery.
- Examples include sendmail, qmail, and postfix.
- Each MTA hop inserts its own data at the start of the email data section.

Edinburgh Napier UNIVERSITY

MDA

- The MDA or Mail Delivery Agent takes email messages from the MTA and delivers it to a particular user or to a MTA.
- Once delivered it is held until an MUA for that user reads the email.
- Examples include mail and procmail for local delivery, and sendmail itself for network delivery.

Edinburgh Napier UNIVERSITY

Example

- Lets consider an example of g.russell@napier.ac.uk delivering an email to me@grussell.org.

Edinburgh Napier UNIVERSITY

1. g.russell MUA on pc236b.napier.ac.uk send the email to the MTA (sendmail) on localhost.
2. The localhost MTA looks up the MX record for grussell.org.
3. The record indicates that email is delivered to grussell.org itself.
4. Sendmail uses its MDA (SMTP) agent to deliver the email to the MTA on grussell.org

Edinburgh Napier UNIVERSITY

5. The MTA on grussell.org looks at the destination user (me) and decides that this is a local user.
6. It uses its local delivery agent MDA (procmail) to put this email into the user's mailbox.
7. The email is stored in /var/spool/mail/me.
8. Next time "me" logs into grussell.org, the email will be waiting for him.

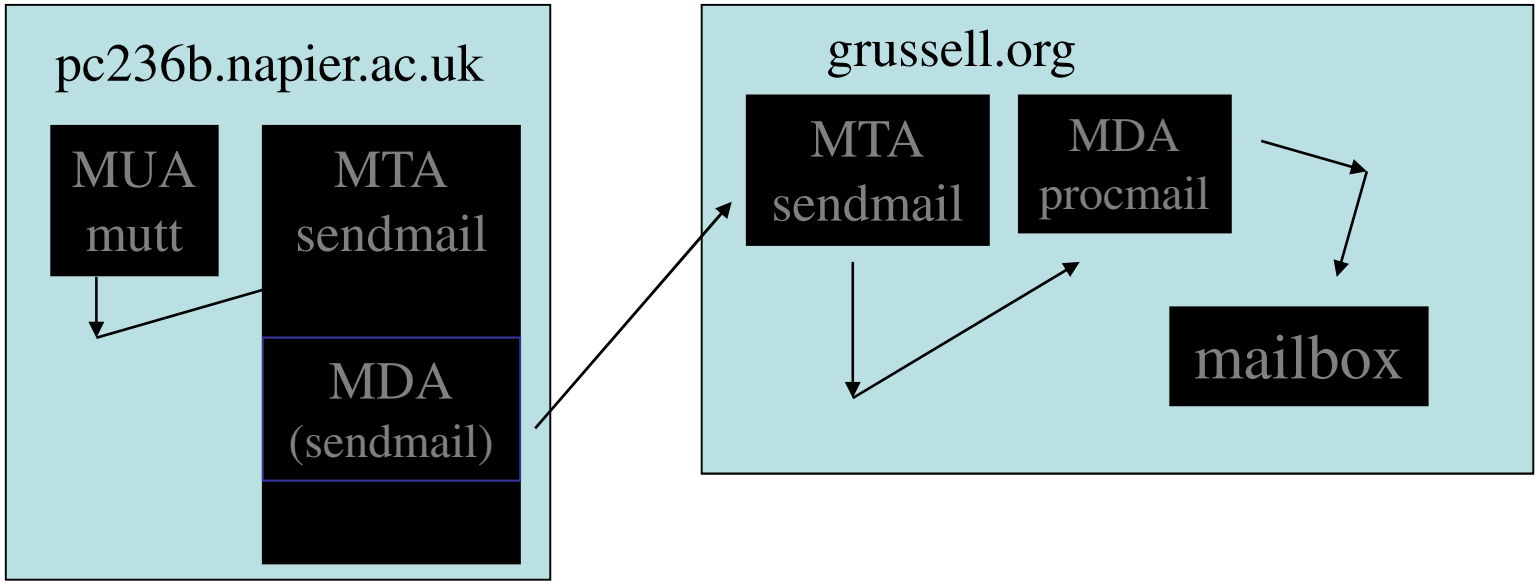
pc236b.napier.ac.uk
MUA
mutt
MTA
sendmail
MDA
(sendmail)
grussell.org
MTA
sendmail
MDA
procmail
mailbox

Edinburgh Napier UNIVERSITY

Sendmail

- Sendmail is the oldest and most common MTA is use today.
- It has many features which are now redundant (like using UUCP to forward email to the destinations using !).
- It is huge and prone to “hack attack”.
- However, it works well, is well understood, universally discussed, and is still popular.

Edinburgh Napier UNIVERSITY

sendmail mqueue processing

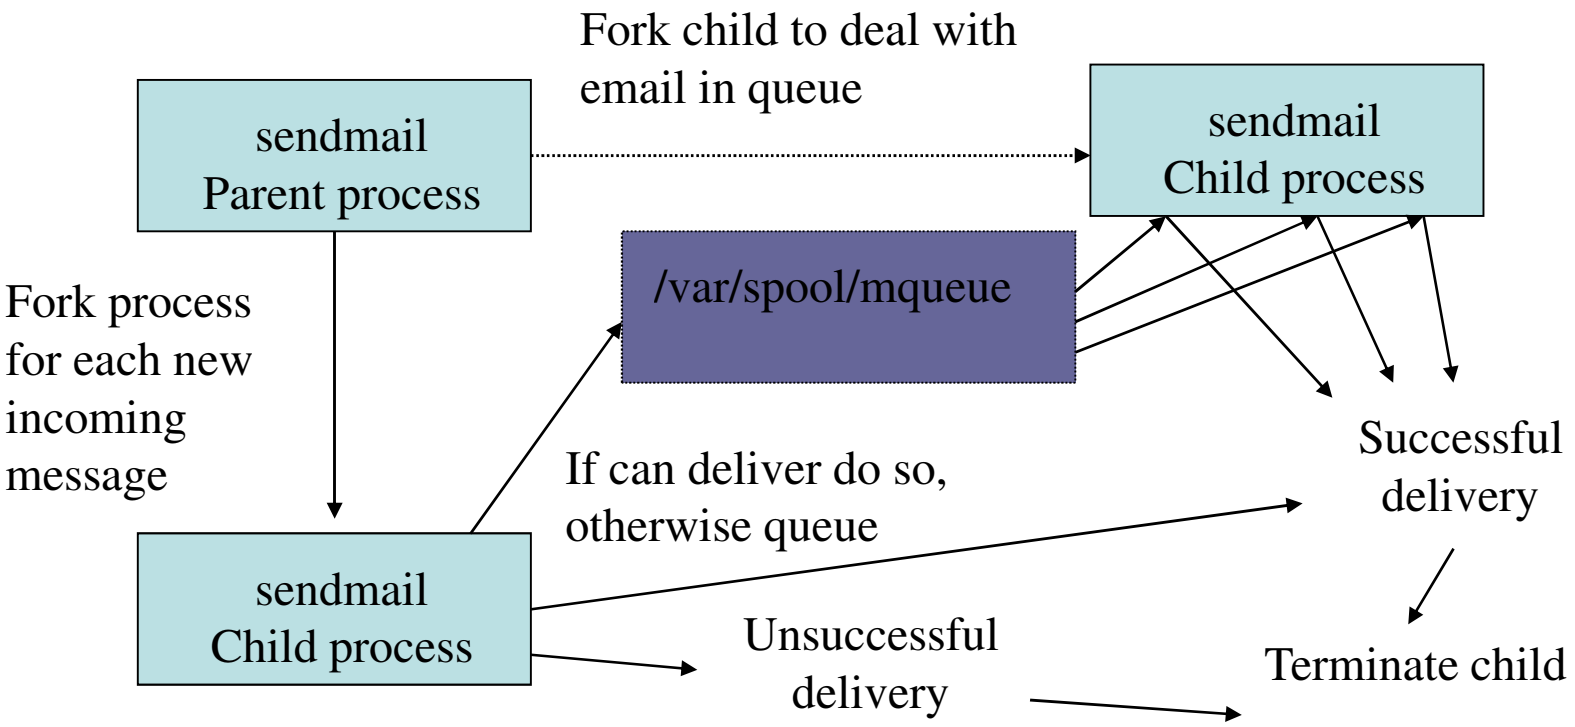

Edinburgh Napier UNIVERSITY

Aliases

- Aliases link the recipient envelope address to a local user or action.
- In sendmail, this is /etc/aliases.
- The file has 1 alias per line, with the alias name, then a :, then the action or user.

Edinburgh Napier UNIVERSITY

Examples

postmaster : root

me : gor

olduser : /dev/null

automail : | /home/gordon/bin/autoregister.pl

devel : gor, me@grussell.org.uk, a.cumming

- Aliases for users, programs (with |) to files (starts with /) or multiple users (separated with ,).

Edinburgh Napier UNIVERSITY

Compiling alias changes

- Sendmail does not use /etc/aliases directly.
- Instead it uses a binary hashed version of the file.
- This is aliases.db or aliases.dbm.
- When you change /etc/aliases you must run newaliases to build the hashed file.

Edinburgh Napier UNIVERSITY

.forward

- Only root can make changes to /etc/aliases.
- Individual users can have simple aliases using .forward in there home directory.
- Every line in this file is treated as an alias for the user.

grussell

andrew

- In this case email for this user would go to both grussell and andrew.

Edinburgh Napier UNIVERSITY

Recursive loops

> cat /home/gordon/.forward

gordon

andrew

- This appears to suggest delivery as normal to gordon, but also a copy of the email to andrew.
- However, for each name the whole alias process is redone, so that (for example) aliases for andrew will also be processed.
- Unfortunately “gordon” will trigger access to his .forward file, find “gordon” again, and again, and again. This would break!

Edinburgh Napier UNIVERSITY

- To tell sendmail not to keep looking things up when they are found in .forward, put a \ character in front of the name. In this example:

\gordon

andrew

- This means “no more aliases on gordon”.

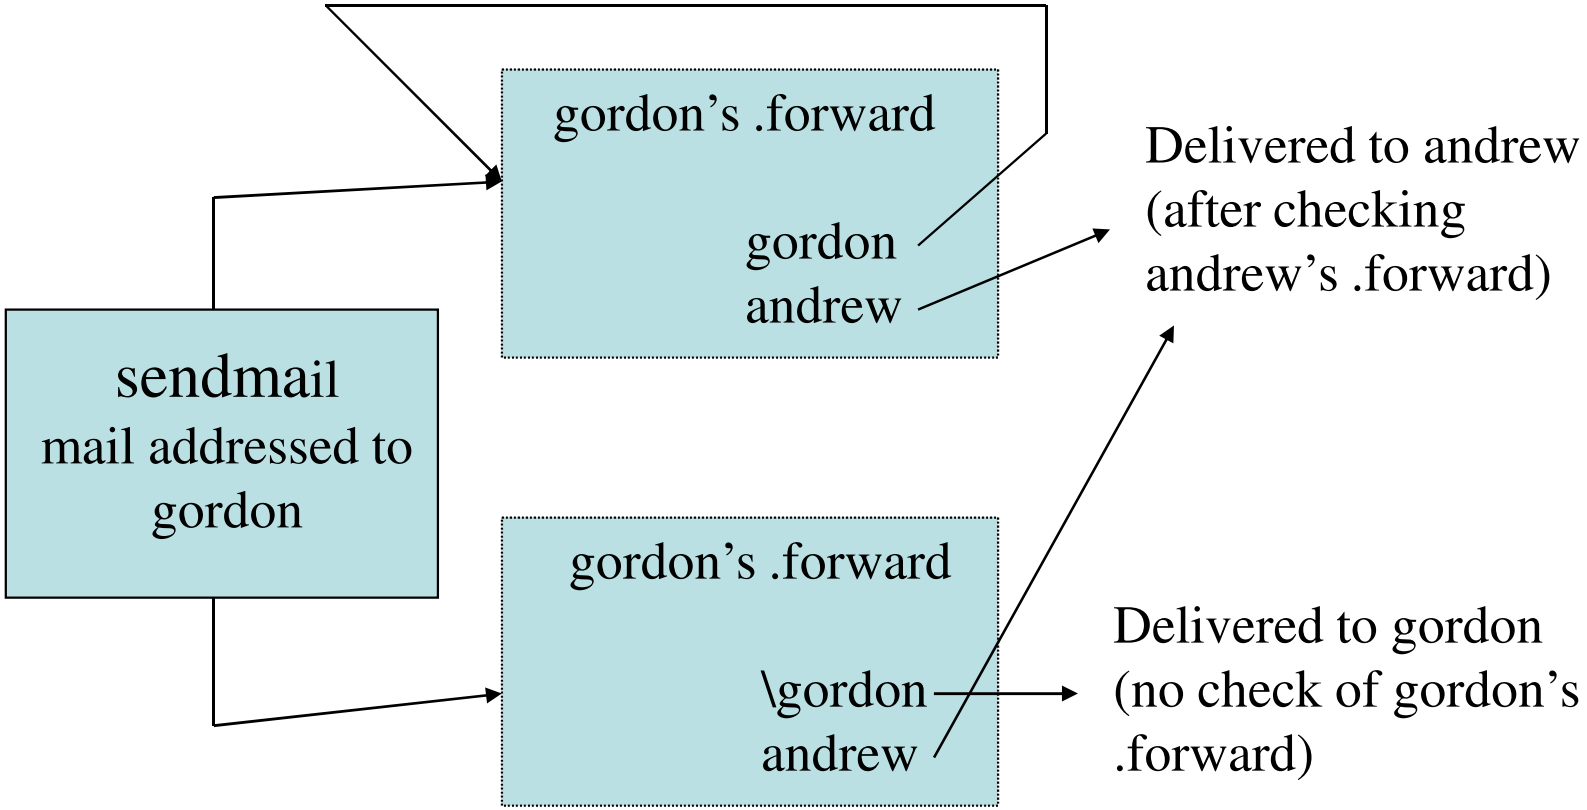
gordon's .forward
gordon
andrew
sendmail
mail addressed to
gordon
gordon's .forward
\gordon
andrew
Delivered to andrew
(after checking
andrew's .forward)
Delivered to gordon
(no check of gordon's
.forward)

Edinburgh Napier UNIVERSITY

Sendmail configuration

- Sendmail conf files are in /etc/mail.
- There is a mix of .cf and .mc files.
- .cf is quickly parsed by sendmail, but for mortals it is hard to understand (but flexible).
- .mc are MACRO level commands which can be used to build (most) .cf files.
- Stick to .mc for all sensible cases!
- Run “make” in this directory to convert from .mc to .cf

Edinburgh Napier UNIVERSITY

Basic structure of sendmail.mc

- OSTYPE()
- define() - 2 parameters: variable name and value.
  define('PROCMAIL_MAILER_PATH','/usr/bin/procmail')
- FEATURE()
- MAILER() – for example
  MAILER(smtp)
  MAILER(procmail)
- dnl – a convenience – like a shell # comment.

Edinburgh Napier UNIVERSITY

What the options are!

- sendmail config is complex!!!
- My sendmail book is >1000 pages long.
- However, the basics are straight-forward.
- Most things you want are either FEATURES, DEFINES, or have their own function name.
- One looks up the effect they want in the documentation, and follows the instructions.
- Only really keen people would try to learn sendmail competely!

Edinburgh Napier UNIVERSITY

Key issues: Masquerade

- There are different types of masquerading in sendmail.
- Masquerading is where the identity of a machine is disguised for some reason.
- For instance, your email server is called “mail.grussell.org”. In a normal setup all your email would appear to come from mail.grussell.org and not just grussell.org.

MASQUERADE_AS('grussell.org')

- Now all outgoing emails are rewritten so that the from address appears as grussell.org.
- It is vital that grussell.org is an A record and not a CNAME.

Edinburgh Napier UNIVERSITY

Relaying

- Relaying is the process of asking one mail server to take your email for delivery somewhere else.
- It can be used by spammers to mass mail junk, but leave the blame on the mail server used.
- It was once switched on by default, which lead to many mail hosts becoming blacklisted.
- /etc/mail/relay-domains contains a list of domains which this server WILL relay.

Edinburgh Napier UNIVERSITY

Virtual Hosts - receiving

- FEATURE(virtualusertable) supports the /etc/mail/virtusertable file.
- Sent emails which match the first column are delivered to that specified in the second column.

me@grussell.org	gordon
me@grussell.org.uk	gordon
jim@grussell.org	error:Sorry he has left
@grussell.org	gradmin

Edinburgh Napier UNIVERSITY

Virtual Hosts - sending

- FEATURE(domaintable) and /etc/mail/domaintable does the virtual host mapping for outgoing email.

gordon	me@grussell.org
andrew	andrew@grussell.org.uk
dbuser	gordon@db.grussell.org

Edinburgh Napier UNIVERSITY

Access

- FEATURE(access_db) and /etc/mail/access allows hosts and domains to be controlled on a per-user basis…

```
lycos.com                DISCARD
grussell.org.uk   RELAY
gordon@sqlzoo.net RELAY
sqlzoo.net               ERROR:5.0.0:550 No way
```

Edinburgh Napier UNIVERSITY

Possible responses.

- OK – local recipients only
- RELAY – forward
- REJECT – no way!
- DISCARD – no way quietly
- ERROR – Send an error back to the sender (RFC 1893 compliant codes).

Edinburgh Napier UNIVERSITY

- There are a number of “black hole” email lists.
- These are lists populated by suspected spammers. Such lists include:
  - http://www.spamhaus.org/
- So to drop people in the sbl list do:

  FEATURE(`dnsbl', `sbl.spamhaus.org', `"550 Mail from " $`'&{client_addr} " refused - see http://www.spamhaus.org/SBL/"")dnl

- See http://www.email-policy.com/Spam-black-lists.htm

Edinburgh Napier UNIVERSITY

Spamhatus

- Spamhatus actually maintains 3 lists:
  - SBL – List of spammers and the like
  - XBL – Machines which seem to be exploited in some way. For instance, have open proxies, or perhaps a worm or virus.
  - PBL – A list of machines which should not really be sending email. For instance, the dynamic ips of customers in an ISP, where they all should be using the ISPs SMTP server for sending email.
- List are only as good as the data used to build them. However, good lists can generally be trusted, so long as you fully understand the implications of using them.

Edinburgh Napier UNIVERSITY

Fighting SPAM

- There are a few research areas where investigators are trying to fight spam.
- One technology being used now is called SPF
  - Sender Policy Framework
- When an email is received suggesting it is from “linuxzoo.net”, the MTA does a DNS lookup on linuxzoo.net, looking for a TXT record.
- If it finds one starting with v=spf1, then some additional checks are performed before the email is processed.

Edinburgh Napier UNIVERSITY

SPF for linuxzoo

- The TXT record supporting SPF in linuxzoo is:
  v=spf1 ip4:146.176.166.1 ip4:146.176.166.15 a ~all
- This indicates that email can only be send by 146.176.166.1, 146.176.166.15, or from an IP which matches the A record for linuxzoo.net.
- If any of those things are true, the mail is processed normally.
- If it is all false, the MTA carries out the “~” action.

Edinburgh Napier UNIVERSITY

Actions

- Really it should have “-all”.
  - This indicates that a failed test results in the email being rejected
- “~all” is a soft fail.
  - This is really for debugging.
  - It indicates that it is probably junk, but I am not brave enough to guarantee it.
  - One day I will change it to “-all”.
  - Until then, it means that if the SPF rules are broken, then delete the email if you are really really being tough.
  - In reality it probably means “take no action”…

Edinburgh Napier
UNIVERSITY

Discussion

Edinburgh Napier UNIVERSITY

Discussion

- A host is acting as an email server for test.com. It is called email.test.com. User jim also uses this server, and wants to get all email send from that machine to be jim@me.com. Any email coming back is to be sent on to jim.
  - Discuss the different aspects of this task.
  - Provide possible virtualhosttable and domaintable entries.

Edinburgh Napier UNIVERSITY

Discussion

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

Edinburgh Napier UNIVERSITY

Question 1

With respect to email delivery, discuss the difference between a header and an envelope.

Edinburgh Napier UNIVERSITY

Question 2

- Provide a virtualusertable for sendmail which supports the following:
  - email to anyone except andrew and jim at hello.com gets delivered to bob locally.
  - email to andrew@hello.com gets delivered locally to andrew.
  - email to jim@hello.com gets delivered to andrew.

Edinburgh Napier UNIVERSITY

Question 3

- When examining an email header, what header elements could be useful in deciding if it is a forged email. Briefly explain your reasoning.