Why ‘user authentication’ is a bad idea

Or: Why weak authentication may be worse than none at all

Nick FitzGerald
Computer Virus Consulting Ltd

Weak authentication < none?

- What is spam?
- SMTP revisited
- Enter SPF, Sender ID, et al.
- Broken before implemented
- Can *spammers* beat it though?
- *Trivially*, and it gets worse…
- So, do we really want to go there?

What this talk is not…

- Dull
- A deeply technical exposition of the piles of truly gnarly brokenness that is SPF and its friends (which alone should prevent any sane folk from considering them)
- Opinionated

What is spam?

- Spam ≡ unsolicted bulk (commercial) Email
- You'd think that may be an important thing to bear in mind if you were developing an 'anti-spam' technology or product…
- …but some seem to have forgotten!

SMTP revisited

- ‘net mail’ one of earliest ARPANET apps
- Finally standardized Aug 1982 RFC 821
- Classic telnet-style text procotol
- Depends on relaying (thus on open relays)
- No authentication
- ‘Open trust’ scheme of ‘early Internet’

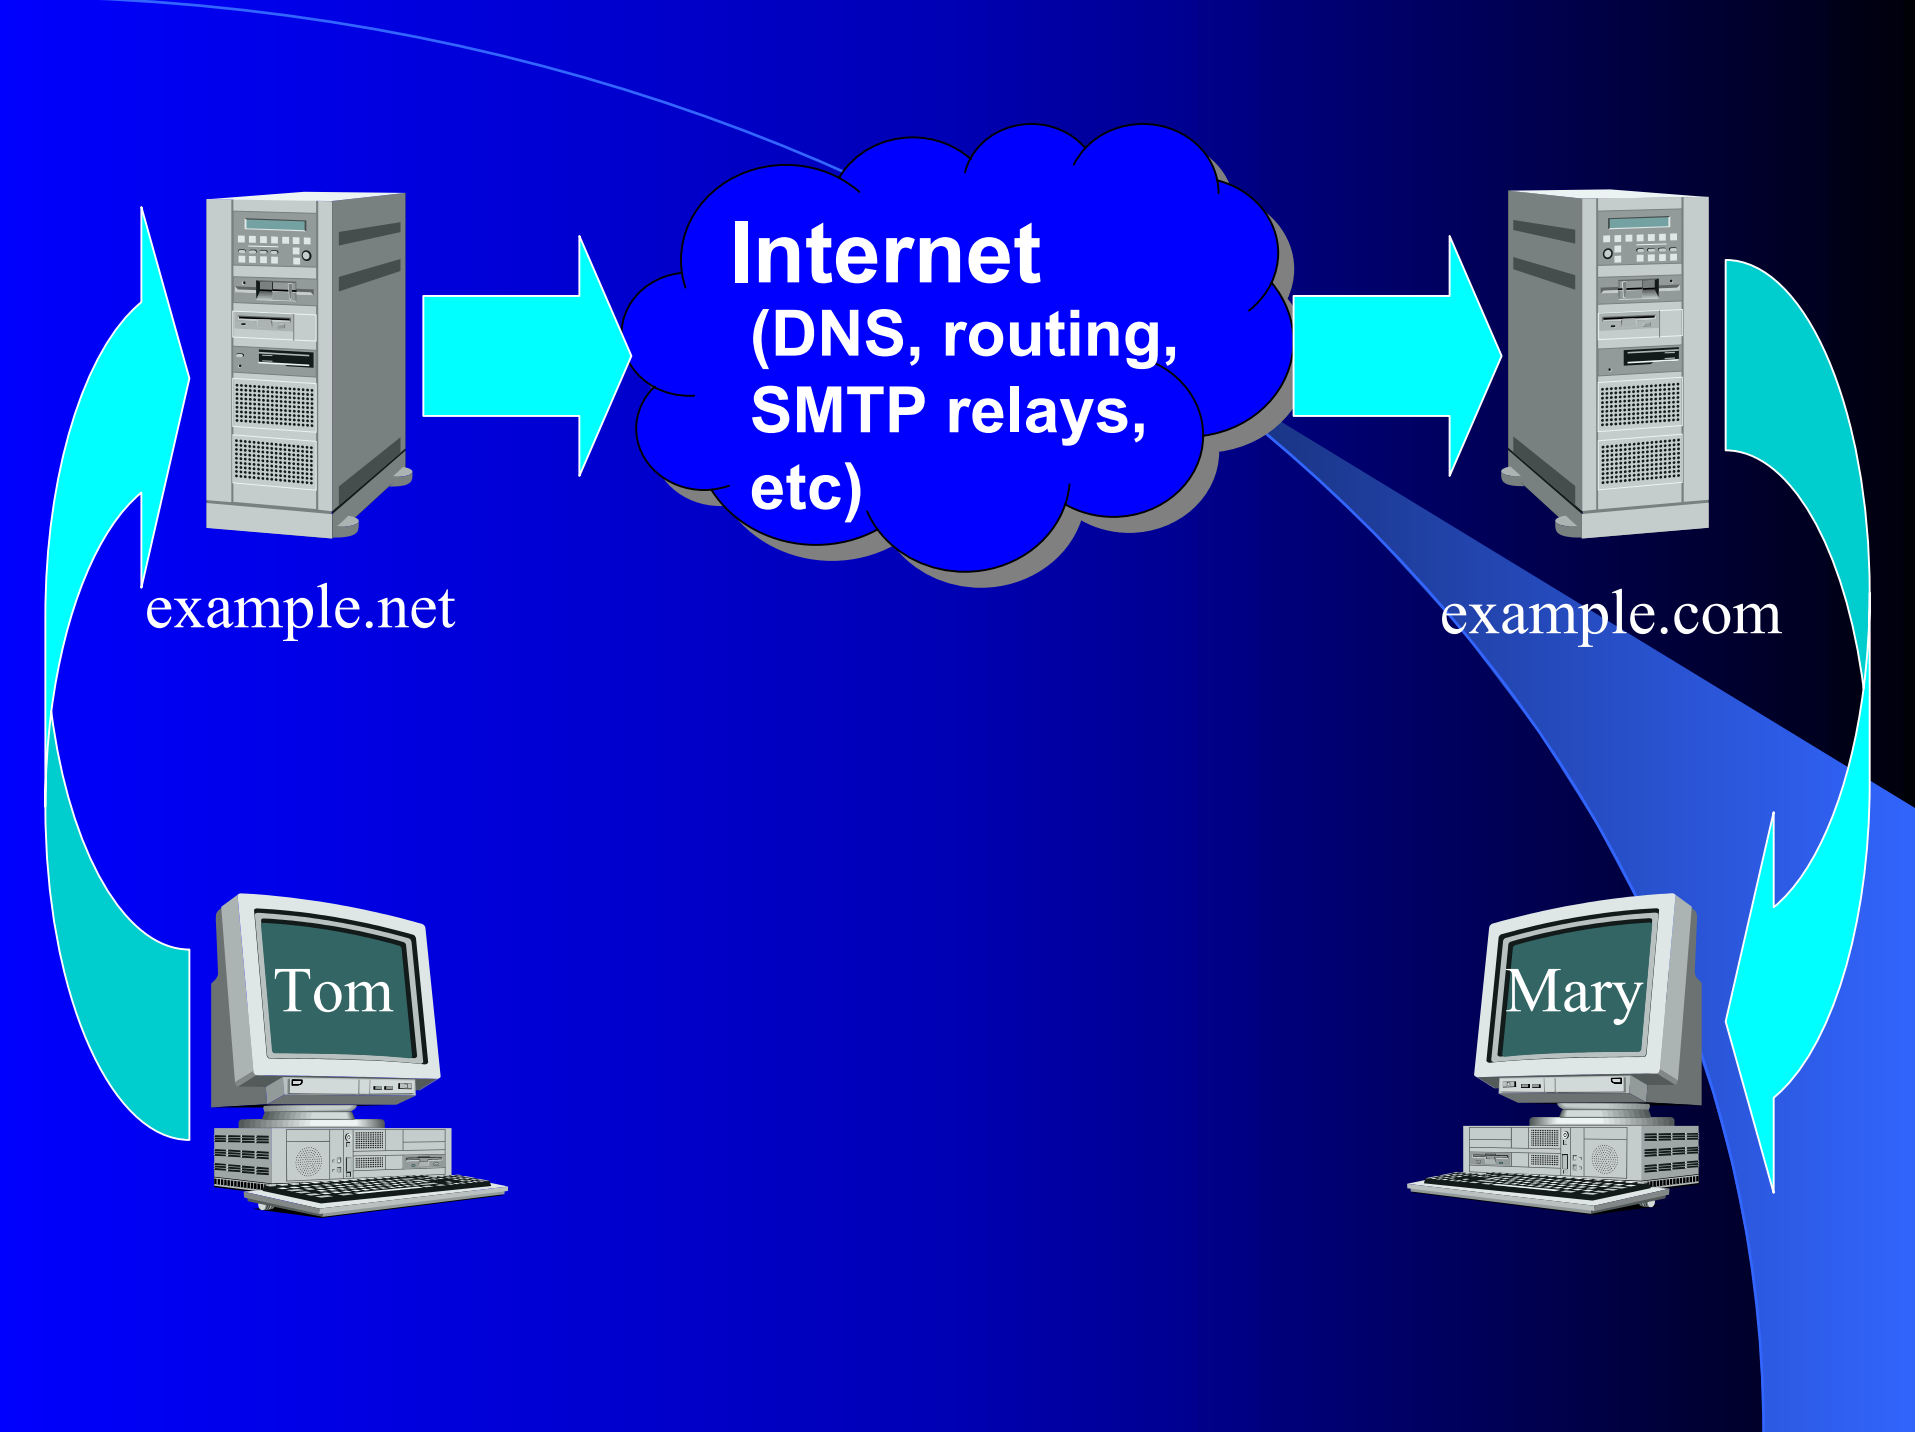
Internet
(DNS, routing, SMTP relays, etc)
example.net
example.com
Tom
Mary

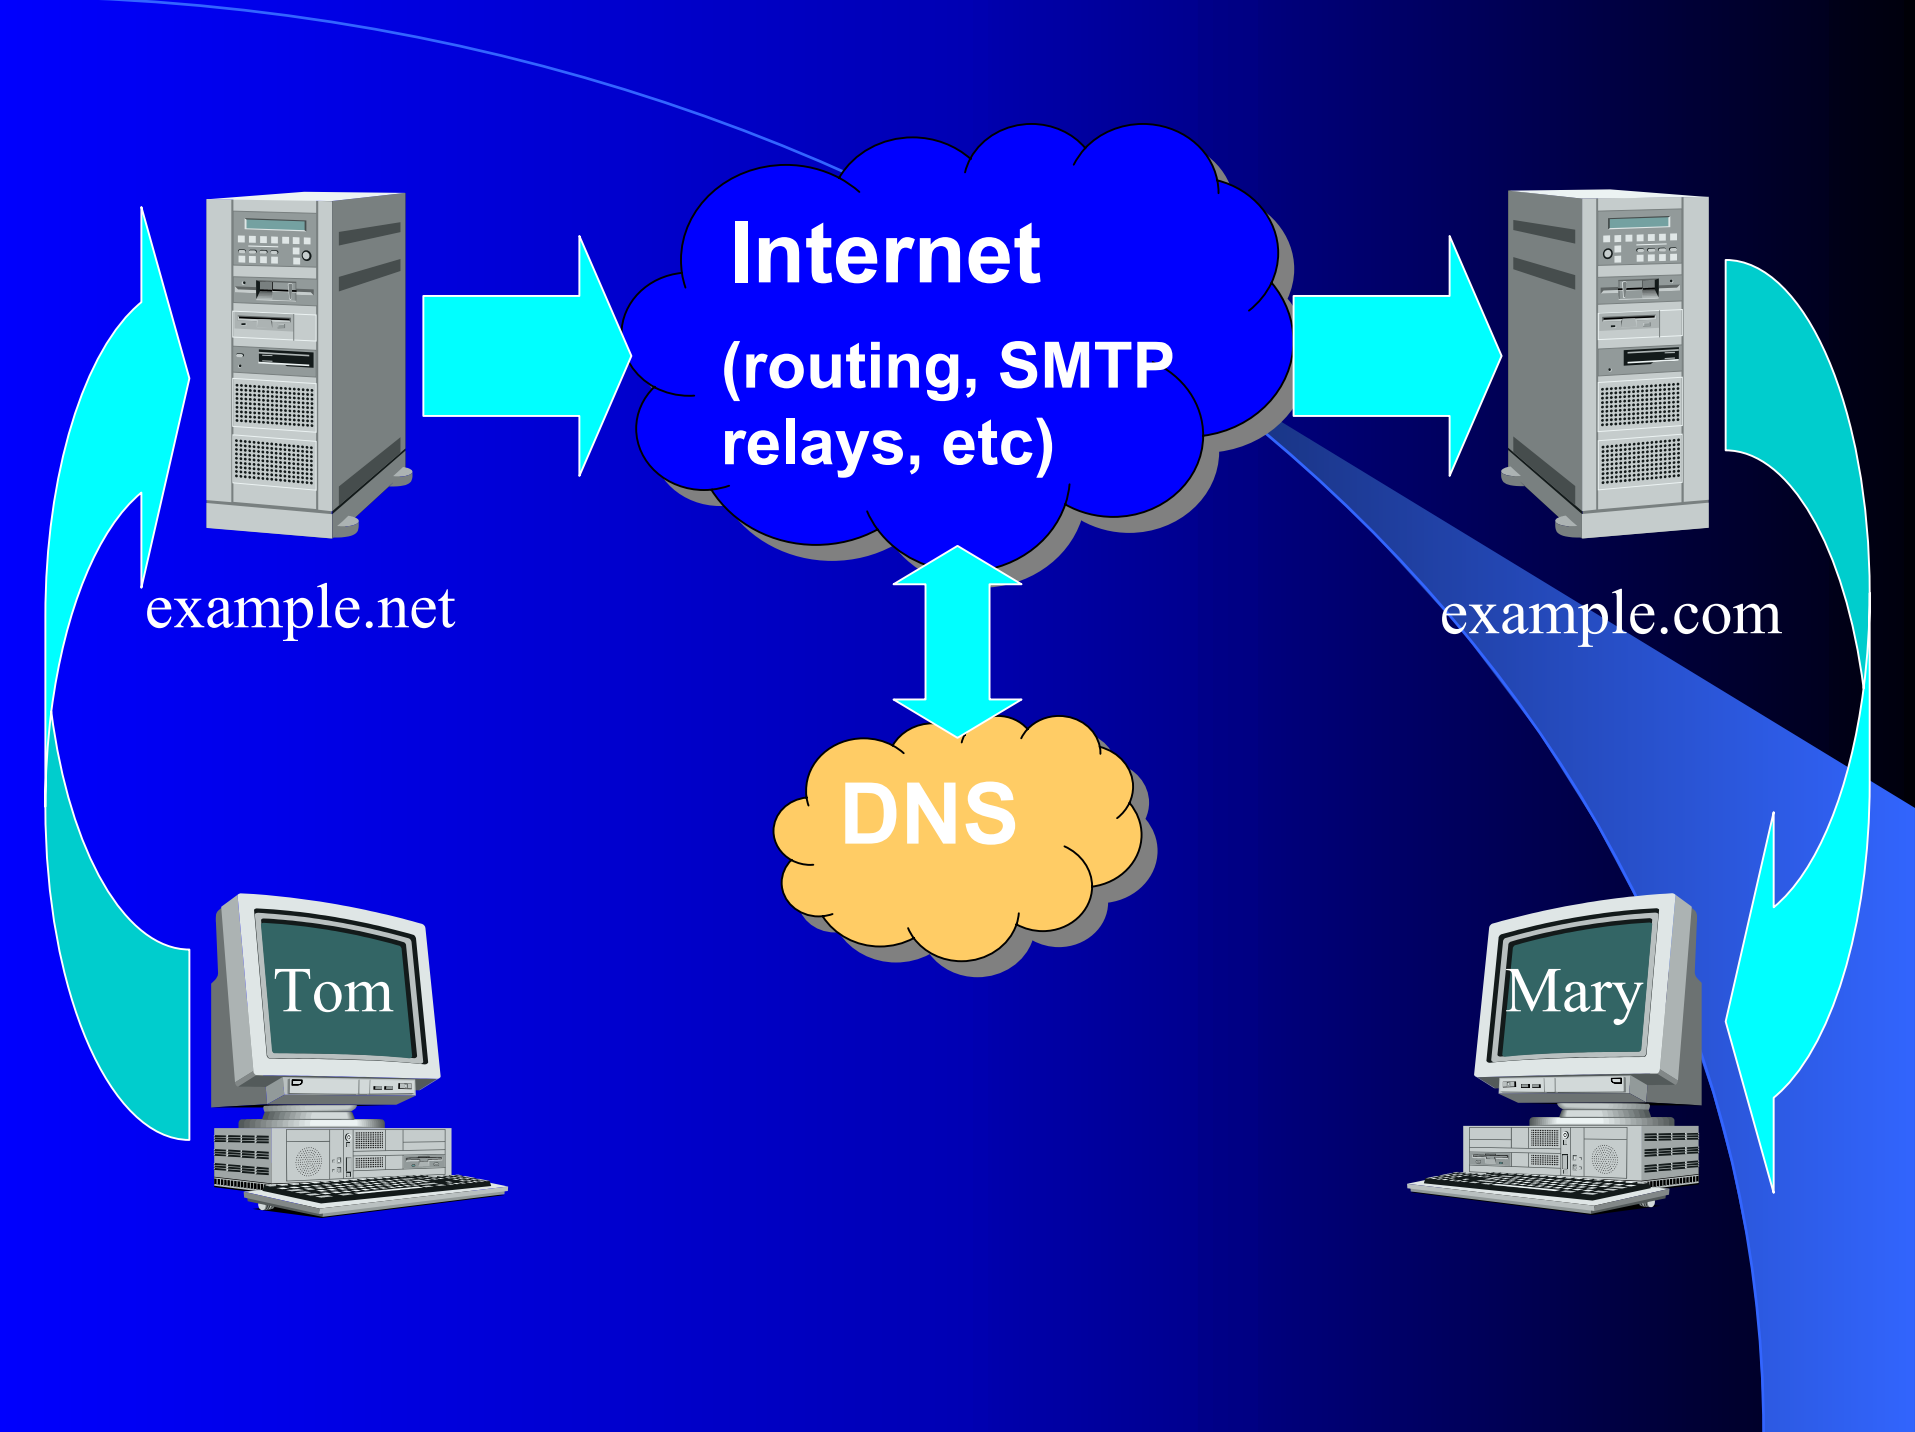
Internet
(routing, SMTP relays, etc)
example.net
example.com
DNS
Tom
Mary

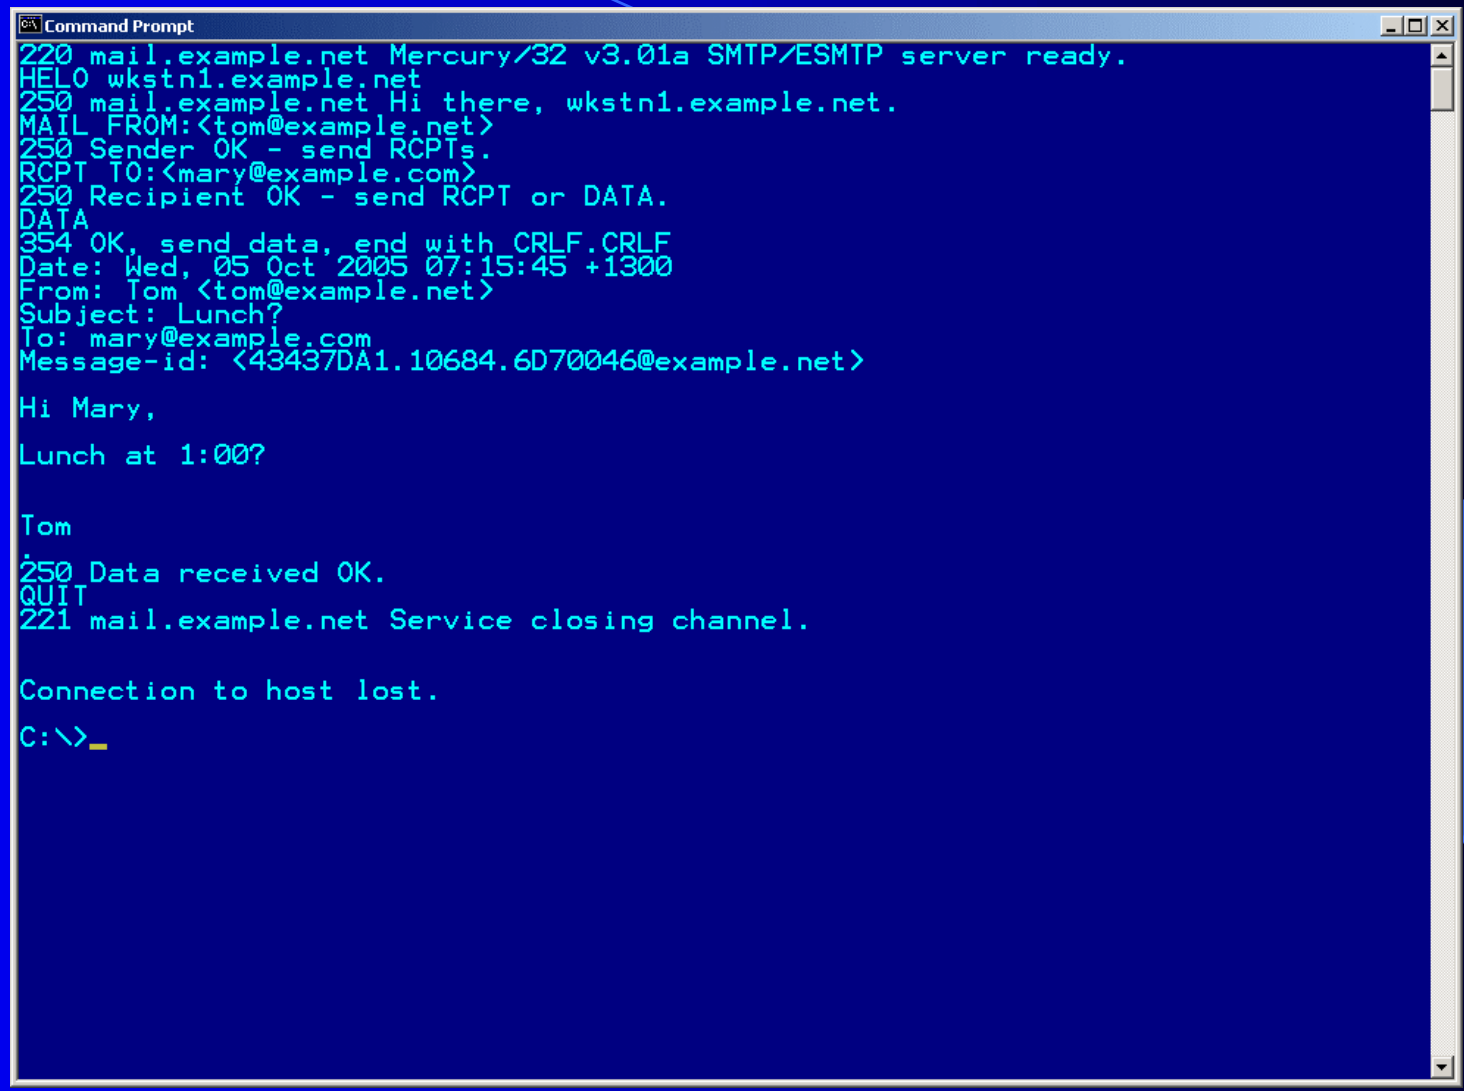
Command Prompt
220 mail.example.net Mercury/32 v3.01a SMTP/ESMTP server ready.
HELO wkstn1.example.net
250 mail.example.net Hi there, wkstn1.example.net.
MAIL FROM:<tom@example.net>
250 Sender OK - send RCPTs.
RCPT TO:<mary@example.com>
250 Recipient OK - send RCPT or DATA.
DATA
354 OK, send data, end with CRLF.CRLF
Date: Wed, 05 Oct 2005 07:15:45 +1300
From: Tom <tom@example.net>
Subject: Lunch?
To: mary@example.com
Message-id: <43437DA1.10684.6D70046@example.net>
Hi Mary,
Lunch at 1:00?
Tom
.
250 Data received OK.
QUIT
221 mail.example.net Service closing channel.
Connection to host lost.
C:\>

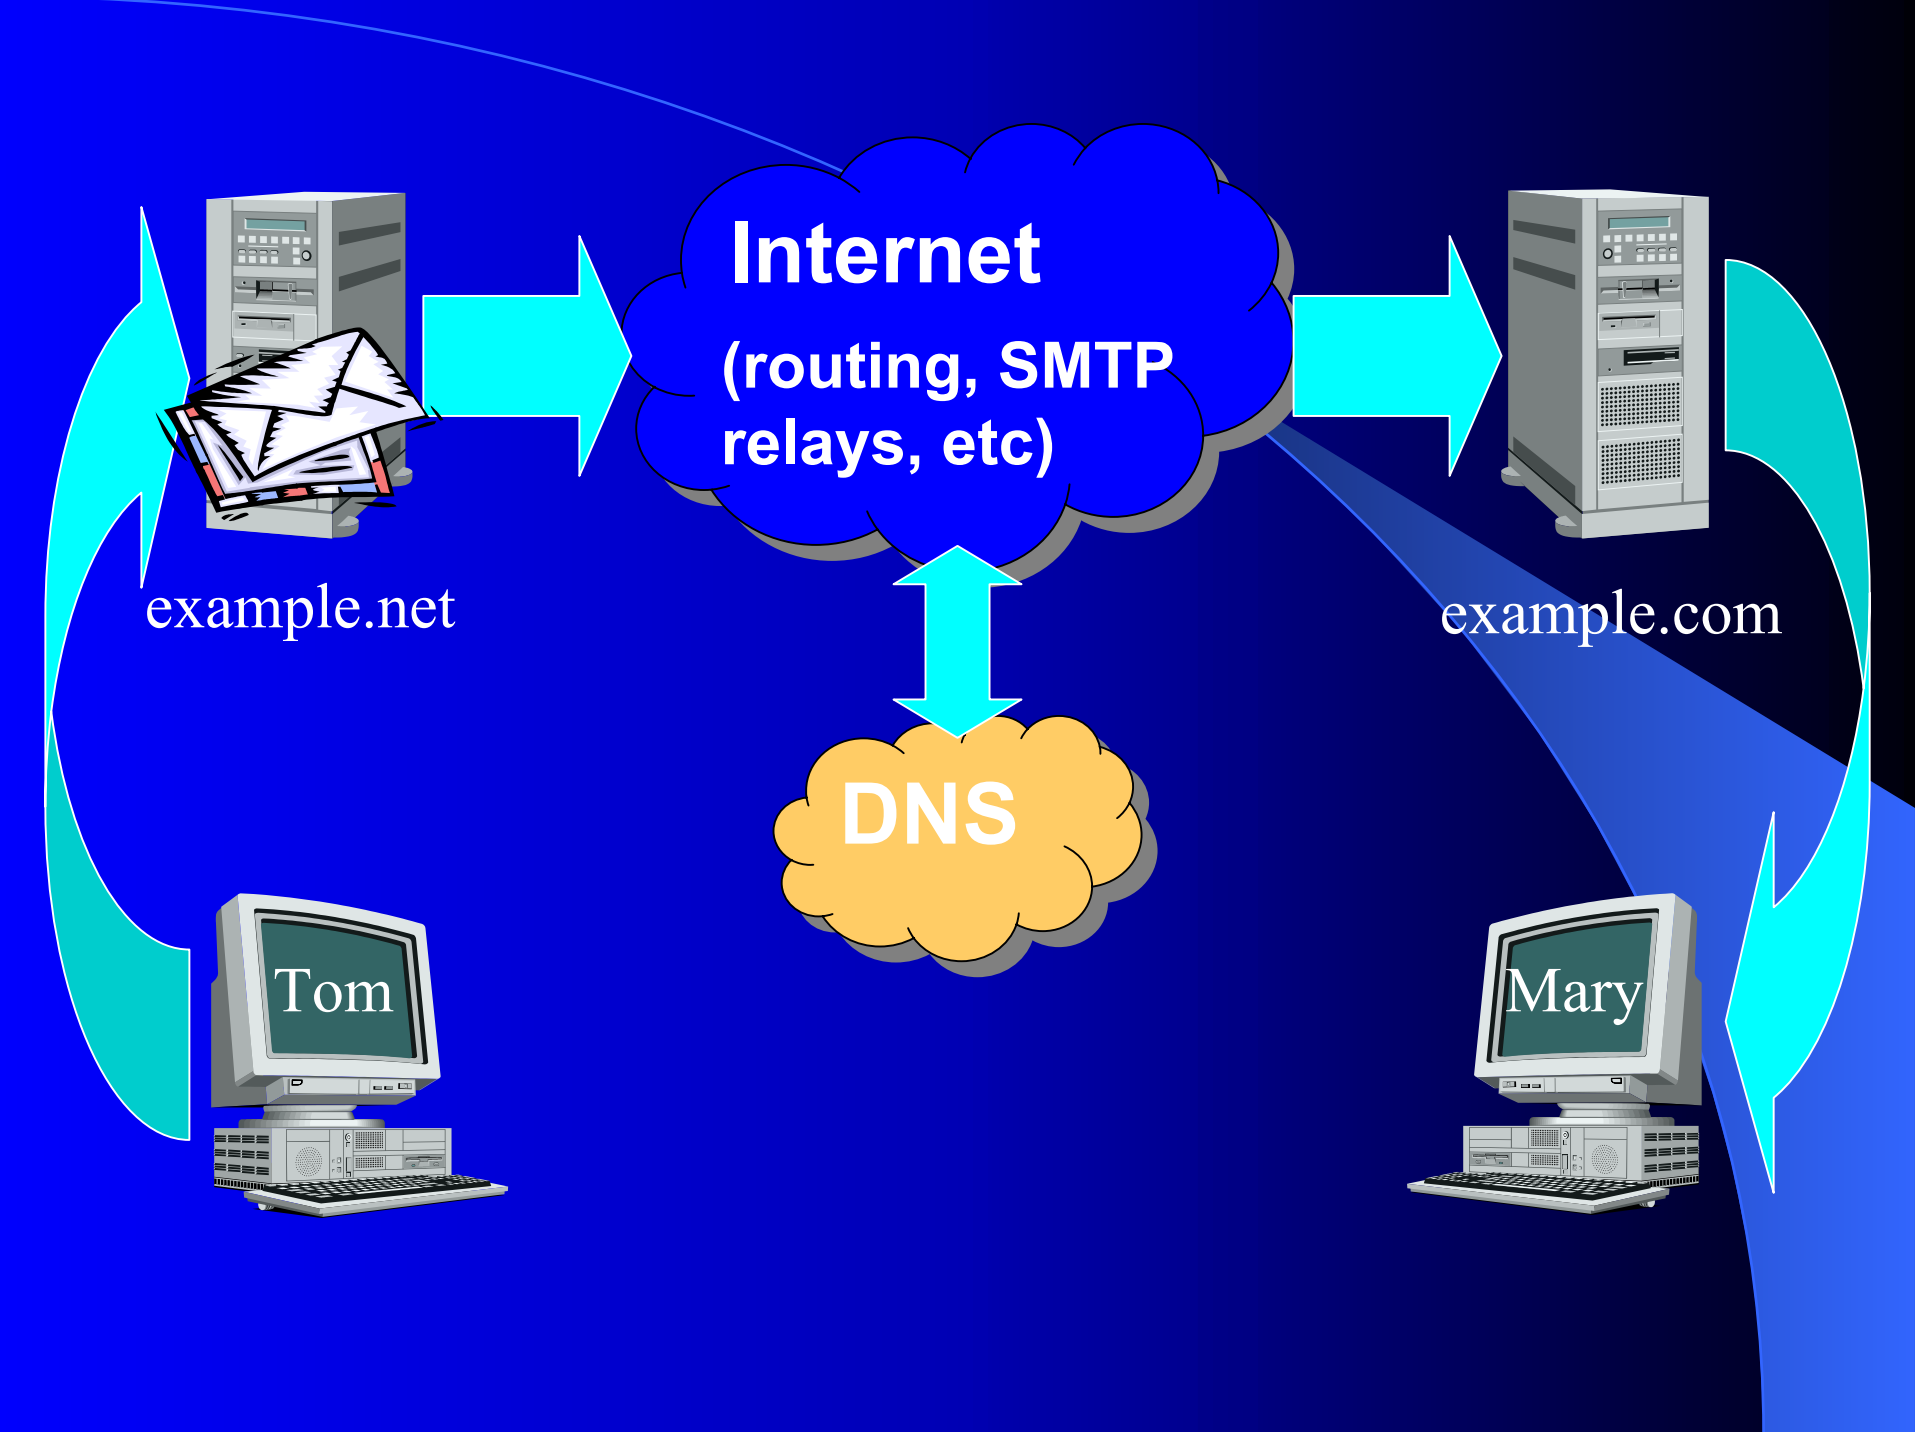
Internet
(routing, SMTP relays, etc)
example.net
example.com
DNS
Tom
Mary

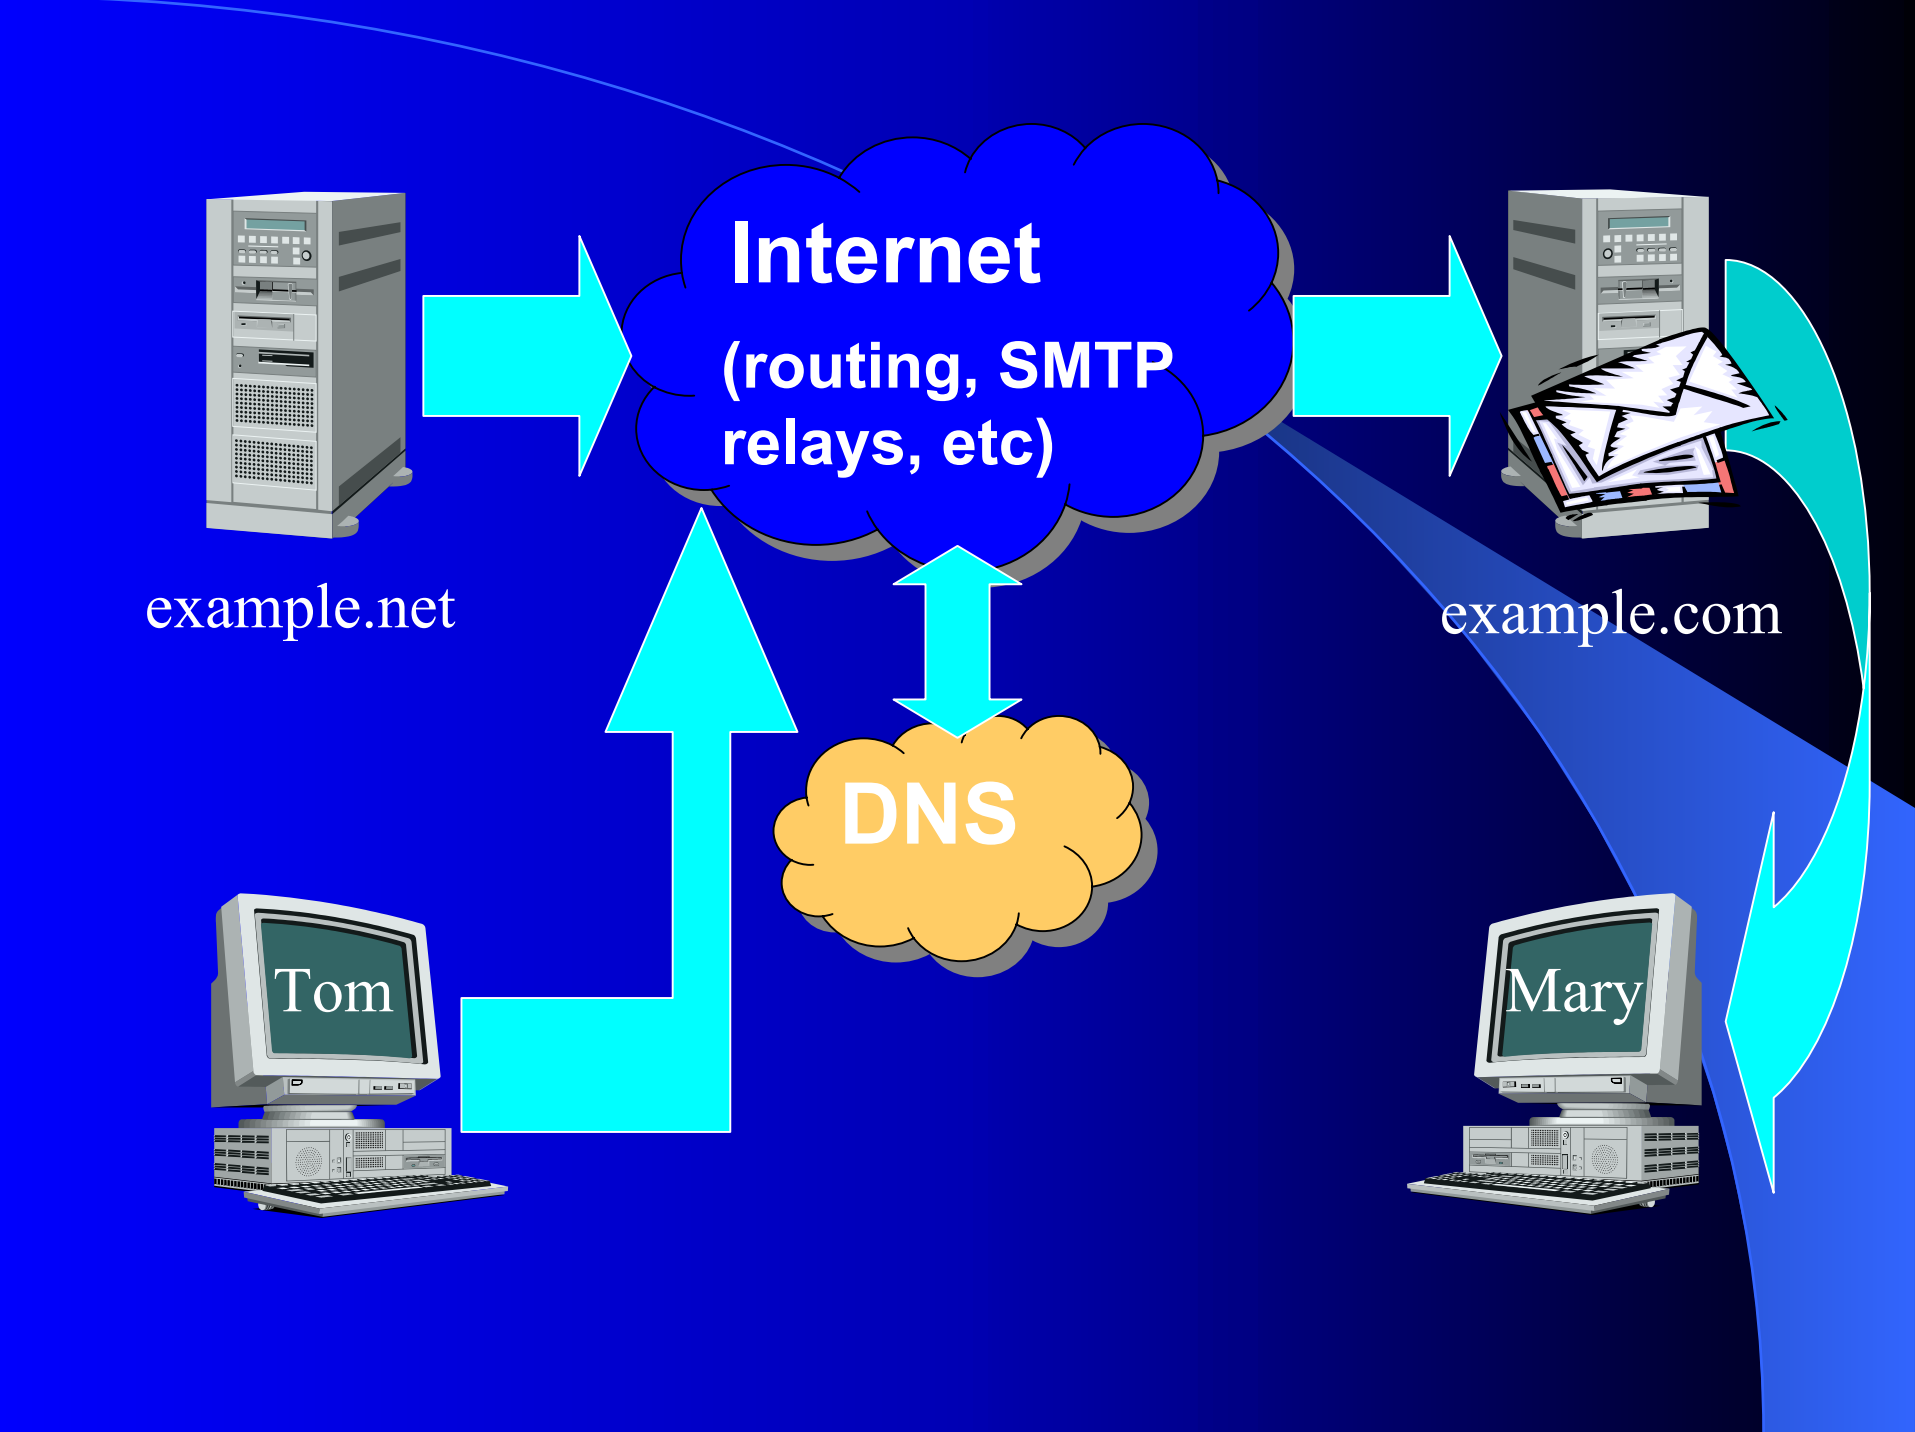
Internet
(routing, SMTP relays, etc)
example.net
example.com
DNS
Tom
Mary

Enter SPF, Sender ID, et al.

- Provide a way to check that sending machines are ‘allowed’ to send Email for the claimed ‘from’ domain
- Recipient SMTP server can do a DNS lookup of claimed ‘from’ domain to see which machines that domain’s admins say can send Email ‘from’ that domain

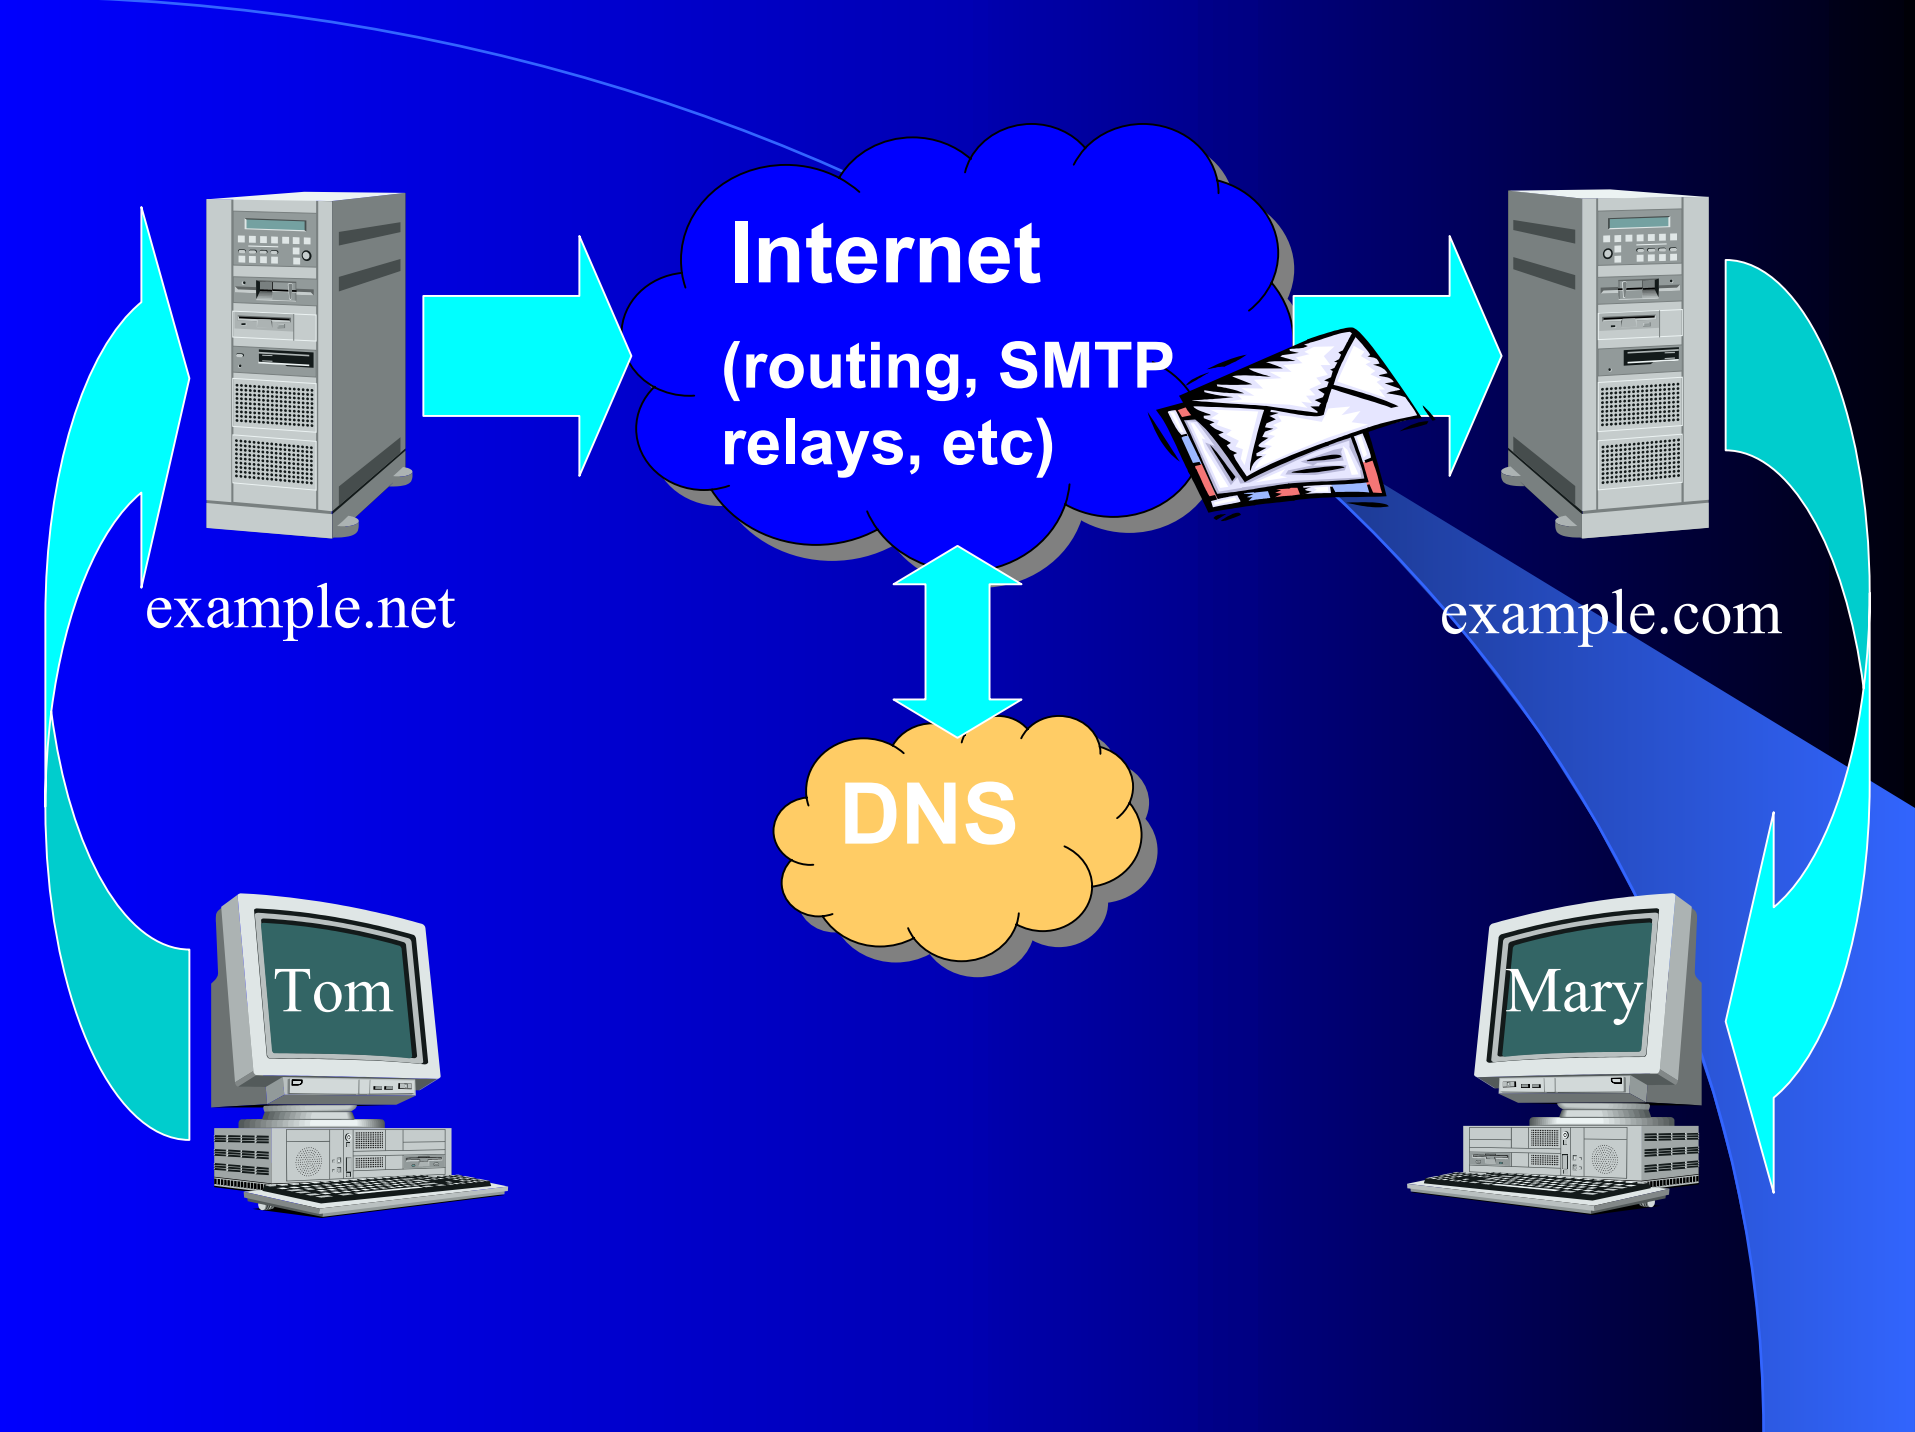
Internet
(routing, SMTP relays, etc)
example.net
example.com
DNS
Tom
Mary

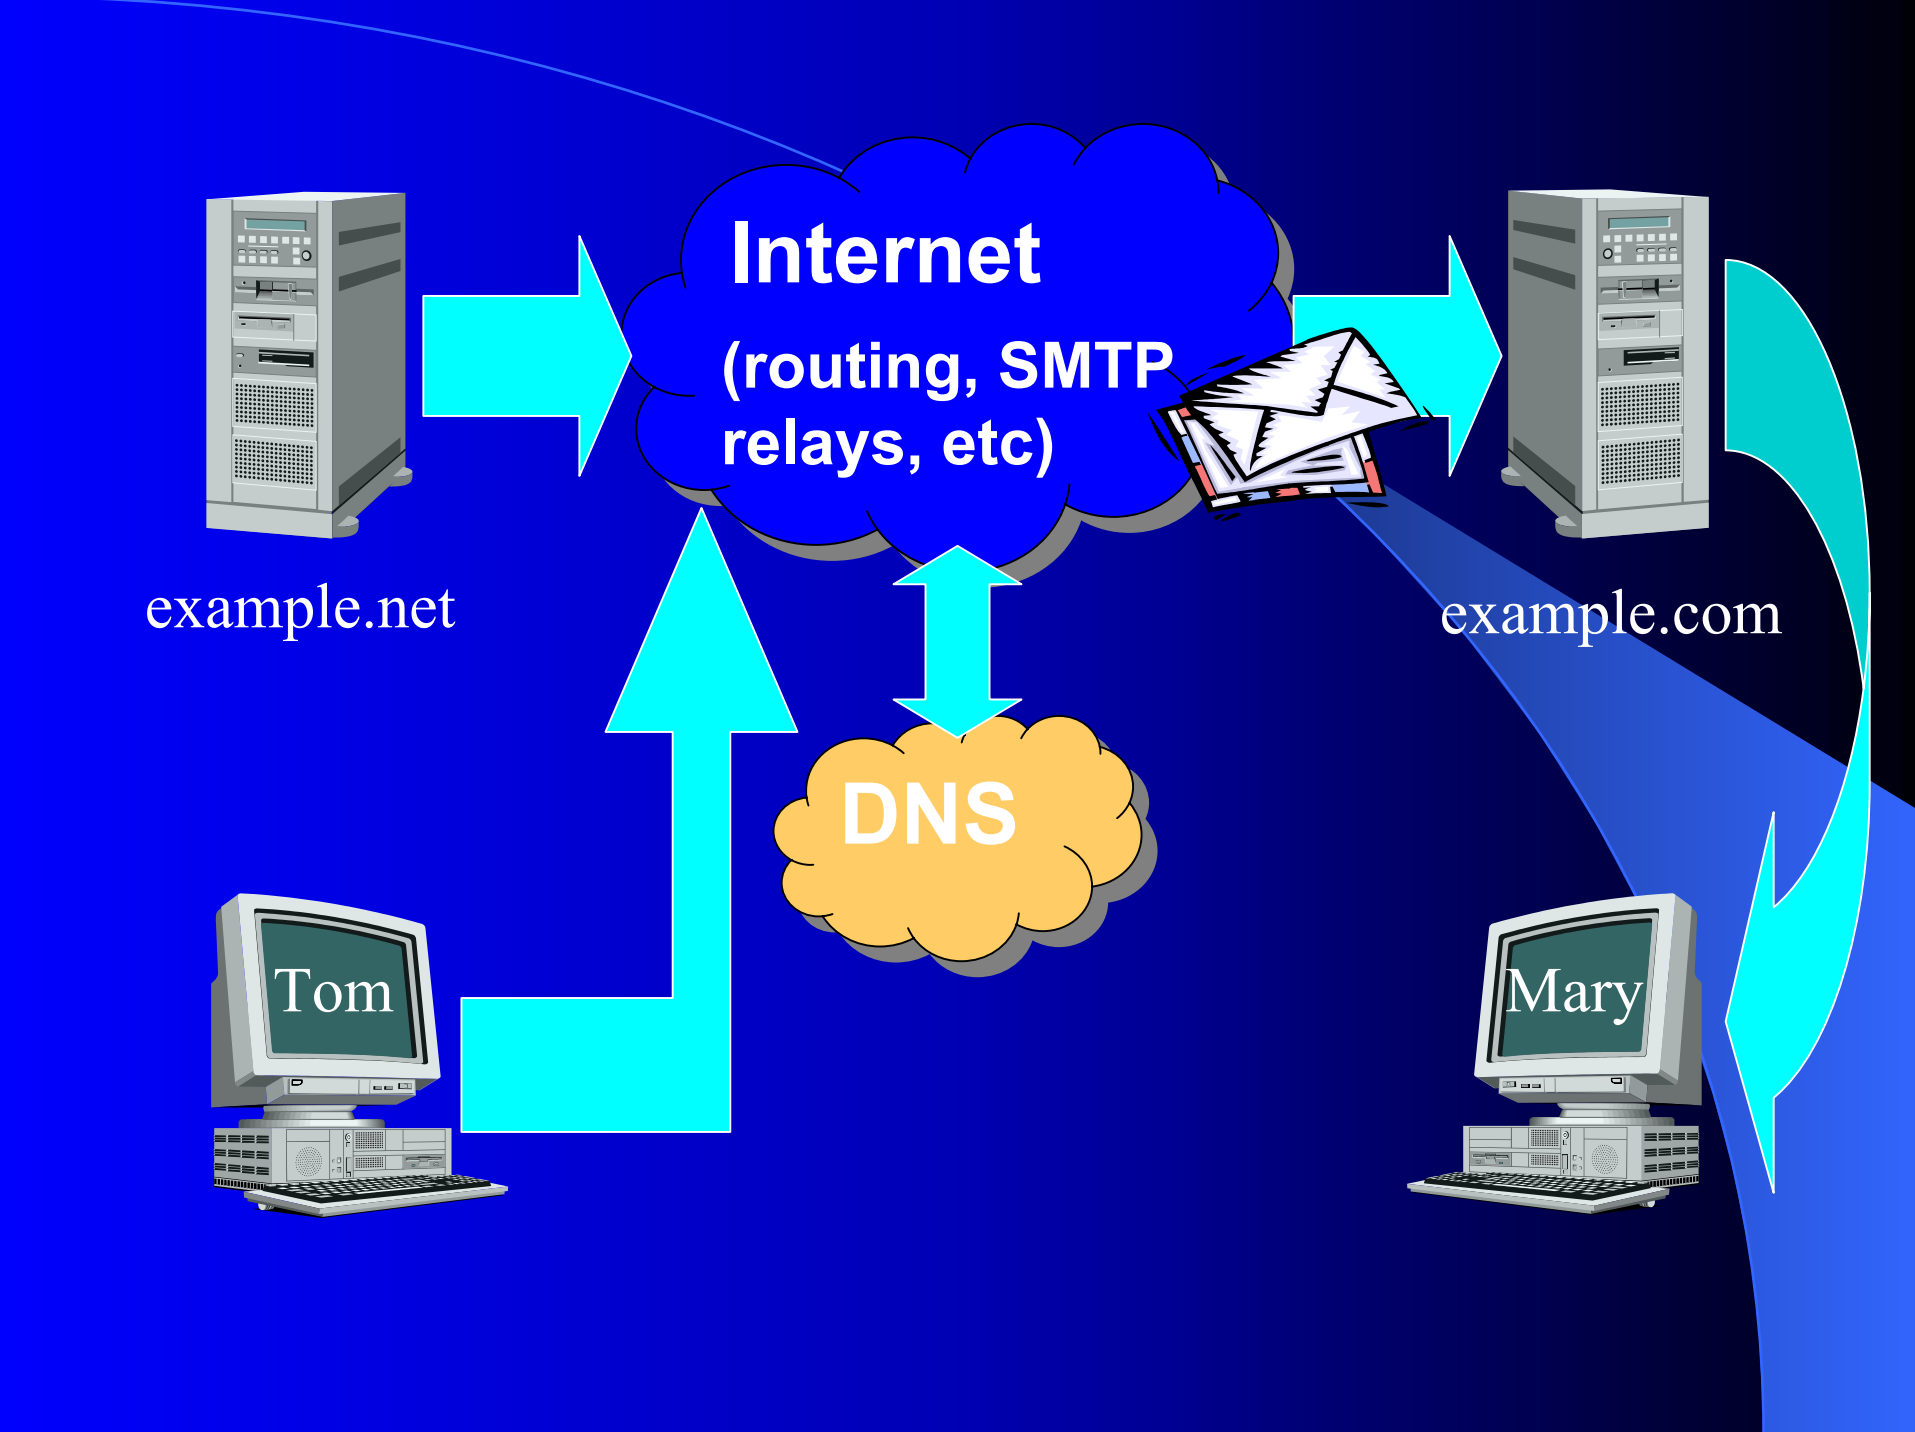
Internet
(routing, SMTP relays, etc)
example.net
example.com
DNS
Tom
Mary

Enter SPF, Sender ID, et al.

- Provide a way to check that sending machines are ‘allowed’ to send Email for the claimed ‘from’ domain
- Recipient SMTP server can do a DNS lookup of claimed ‘from’ domain to see which machines that domain’s admins say can send Email ‘from’ that domain
- Widely referred to as ‘user authentication’ and other such nonsense

Nonsense?

- Where is the ‘authentication’ done?
- At the network connection endpoint level

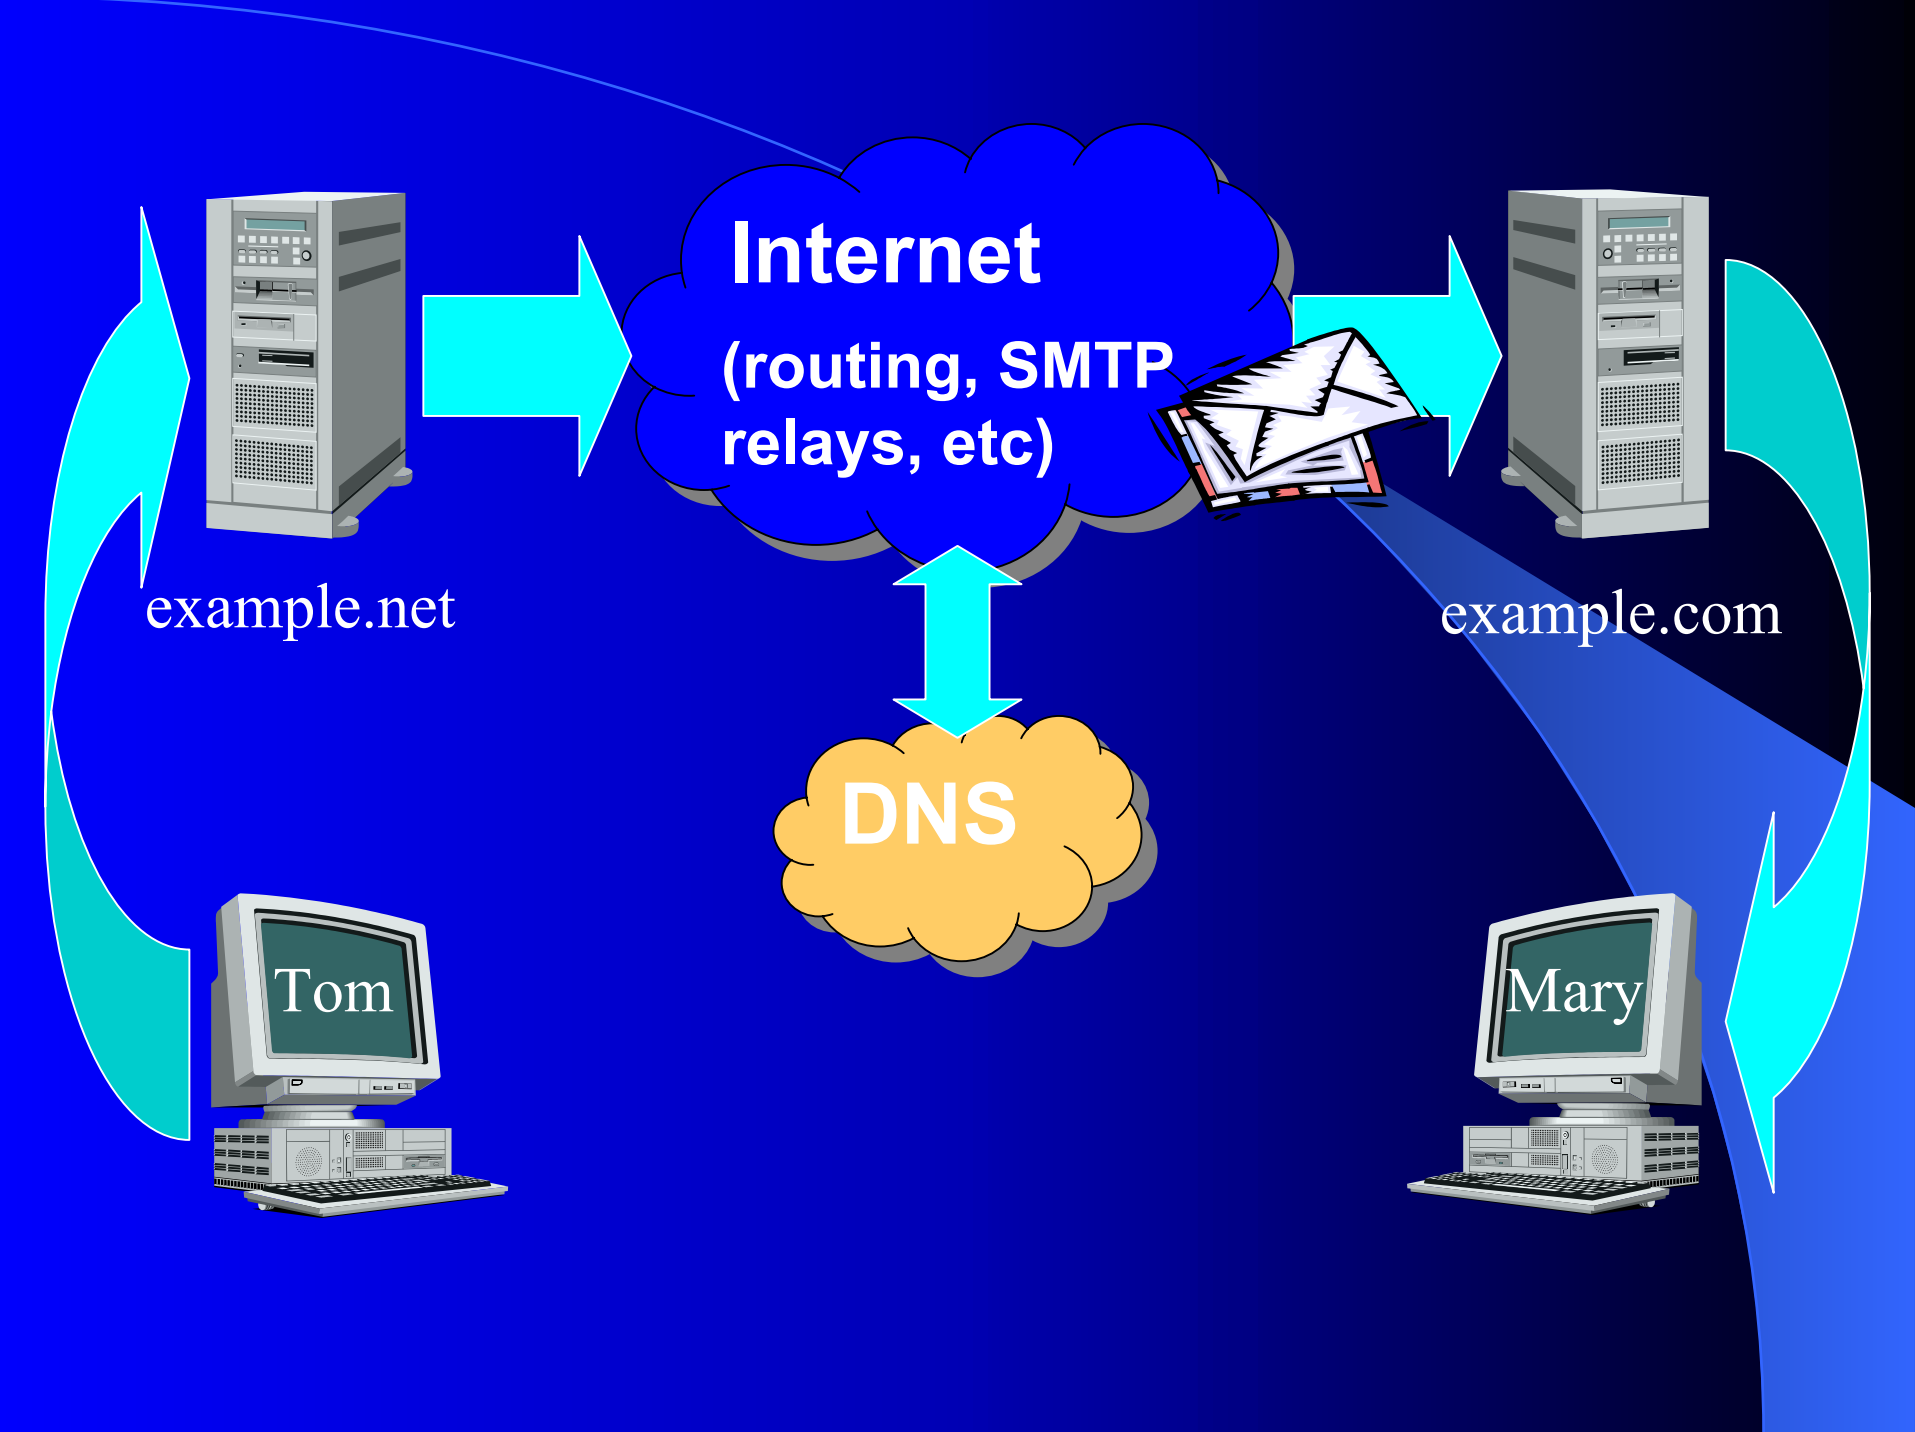
Internet
(routing, SMTP relays, etc)
example.net
example.com
DNS
Tom
Mary

Nonsense?

- Where is the ‘authentication’ done?
- At the network connection endpoint level
- Same as ‘Caller ID’ in the phone network…
- …and that is more correctly known as Caller *Line* Identification (CLI)
- There is no ‘user’-anything involved here
- *Any* process sending via an ‘SPF approved’ server can send SPF-compliant messages

Broken before implemented

- Given
  - Spam ≡ unsolicted bulk (commercial) Email
  - *Any* process sending via an ‘SPF approved’ server can send SPF-compliant messages
- Knowing a message arrived SPF-compliantly tells us nothing about
  - Its **actual** sender
  - Its spamminess
- Result ≡ broken before implemented

Can spammers beat it though?

- Trivially
- They already have large botnets
- ~80% of spam from compromised PCs running:
  - SMTP relay
  - Dedicated spam-bot
- Current spam-bots don't directly beat SPF…
- …but it is trivial to add a few lines of code to them to 'fix' that

Trivially?

- Yep. Recall, on a botted machine, the first immutable security law already applies:
  - Once a bad guy runs his program on your computer, it's not your computer anymore.
- A spam-bot could easily:
  - [elided to not help the bad guys]
  - "
  - "
  - "

But, it gets worse…

- Spam-bots could easily be modified to:
  - [elided to not help the bad guys]
  - "
  - heaps of other gnarly stuff I (and I'm sure the spam-bot writers) would think of had I spent more than five seconds on it

…and worse…

- For SPF to be ‘useful’, it needs a (substantial) critical mass
- At a substantial cost to those choosing to adopt
- Before that critical mass is reached you *can* (will) see an improvement in your spam blocking because SPF will *incidentally* block spam because of common, but non-essential, features of today’s spam…

…and worse…

- …but well before that critical mass is attained, the spammers will start to feel the pinch…
- …which means they'll respond
- They'll talk to their bot developers, work out 'fixes' something like those I have suggested and pay to have these changes implemented…

…and worse…

- A few days later they will push out the next update to their bots and we'll see most bot-sent spam become fully, irrevocably and forever SPF-compliant
- Any further tightening of the screws and they'll move to only spamming SPF-compliantly from within each bot-hosting network

So, do we really want to go there?

NO!!!!!!!!!!!!!!!

Questions?

Nick FitzGerald
nick@virus-l.demon.co.uk