

## HOW DOES EMAIL WORK?

\*Gordon Woolf sees it as surprising that email works as well as it does, but we should be aware of some of the potential problems

For a while I thought of dropping the word "how" and asking "does email work?"

Perhaps we should be surprised that it works as well as it does, but I was prompted to write this when a friend said she had not called on a relative because "he won't even answer my emails".

I asked what to me seemed the obvious question: "Are you sure he got them?"

We tend to take email for granted, and while it works well perhaps 99.9 per cent there should never be an assumption that an email gets there unless you receive a reply.

It worries me that people are increasingly transferring non-urgent but critical mail to the Internet. You can see the problems with real mail: did you in fact post it, did it have the right address or postcode or country, did it get put in the right mail box in the right street, did it get taken out of the box by the right person, did it get put on the right desk or in-tray, did it fall into the waste basket, did the right person read it, did they understand it, do they want to do anything as a result?

Be amazed -- but no maze

If all those things can go wrong with normal mail we should be amazed that things usually arrive, and within a day or two.

We should be even more amazed that email arrives and often on the other side of the world within a minute or two.

However, it is not the somewhat disorganised maze of the early days. Since the mid 1990s, worldwide, it has been a very organized network owned by big businesses and some governments such as, in Australia by the Singapore and (seemingly reluctantly) the Australian government.

Most ISPs, including large ones, use the "backbones" provided by major communications providers -- the telcos of the world -- and you will generally find those datacenters where the computers of service providers are installed, are at the major junctions in telephone cables... just as cities developed at the junctions of trade routes. There are around 8 to 12 tier-one internet service providers in the world, and they are big business with names you'll know like AT&T, MCI, NTT/Verio, Sprint, QWest and Savvis (which was the US Cable & Wireless). While it is common for them to work on a "peering" basis in that no money changes hands for using each others' networks, disputes are not unknown. One in October this year between two companies which may or may not quite qualify as being in "tier one" caused problems for millions of direct and email connections when they refused to deal with each other for a few days.

In Australia the backbone providers are Telstra (by far the biggest) and such as Optus, Connect.com, UUNET and the revitalised universities network now contracted to Optus. Some of the smaller ISPs have direct peering arrangements with other smaller ISPs to route traffic to each other without going through the major networks.

At a much more local level, internet service providers can depend on being near the major cable routes. Several years ago a small computer retailer in a Victorian country town wondered why he continued to be approached by service providers large and small who wanted to site a computer in his back office. The reason became obvious when one looked out of his back door. At the end of the laneway, perhaps 10 metres away, was the back gate to the main telephone exchange for the region, hidden away behind what used to be a post office.

Where does it go?

Who provides your email service? The setup between companies at all levels can be confusing. To take an actual example, you could get an email address from a small web host based in Australia but find that the company which provides the technical services is in California, USA. The computers are actually installed in a datacentre (or should that be datacenter) in Michigan (with a backup on the US east coast). That datacenter is owned by one of the companies which do nothing but provide secure space for racks of computers with connections to the Internet backbone (the tier-one businesses) in multiple ways. They have technical staff who deal with hardware faults and system software, but everything else is controlled by the business in Australia. That is how international we are becoming.

Even if you get your ISP service from a big supplier like Optus (owned by SingTel and therefore ultimately majority owned by the Singapore Government) you will most likely be connected via Telstra's wires to the local exchange but at some point, either there or at the regional exchange you will start to go on the Optus network and to one of their datacentres.

Between major cities you may well be on one or other of the major carrier's cables, but remember when a ditchdigger cut into the connection at Wangaratta a few years ago, it was found that the two major cables from Melbourne to Sydney were at that point in the same ditch. There were backups, but not sufficient to cope seamlessly with the two major cables going out at the same time.

When power went out in a large part of the USA due to some faulty planning on dealing with overloads, it was found that two of the seven computers which provided all the web and email address routing information for the world were not only in the same city, but in the same building. That is no longer true... there are more of them and exactly where they are is no longer common knowledge.

Typically, the datacentre my own website uses has direct links to three of the tier-one providers and via them has 4500 networks just an extra hop away. That's by no means a great service, but it is affordable.

However, we should not forget that the most common reason why an email does not get to the person it is intended for is that the sender typed the wrong address. It can be as simple as forgetting to press the shift key to get the @ sign (well everything else can be typed without the shift!) Most email programs will tell you if you attempt to send an email without the @, but not all.

What should happen if a mistyped address gets sent is that it should bounce back, but it is always possible that your mistyping produced a real address, the recipient knows not what you are writing about, regards it as spam and deletes it.

The postmaster

One way of checking on addresses is to write to the address which, according to the email standards, should exist for every domain. That is "postmaster@domain.com". Unfortunately some domains are so large that they can't help (put Bigpond and Yahoo in that category, and many smaller ones are set up either without a "postmaster" account or have it diverting automatically to that place called ":null" or ":blackhole" from where nothing ever returns.

The other annoying real life bounces are those which tell you the mailbox is "over quota", meaning it is probably full of spam and that the user has gone on annual holidays without thinking what may happen to email. Some systems usefully tell you that a 1k message would be accepted and let the account accept a message which could say something like "Hey, Joe, let me know when you clear your email box. I want to send you something." Do switch off that nice graphic signature or you'll just get another bounce.

And it might not be spam which is the reason. If you send an email with those pictures of the kids which you did not check the sizes of, you could be the cause of others getting the above message. The most common email inbox size limit is 10MB, but there are many which have a 2MB limit or less. Of course GMail inboxes are huge -- you'll never fill those, but, on the other hand it is likely that items sent via GMail will never be deleted - ever, even if you send

them to the trash.

You may also be sent a bounce which says that mail could not be delivered for a specific period, maybe a few hours or maybe a few days, and that the system will keep trying for a period of maybe three or five days. These do mean what they say, but if a fault has lasted several days already it is likely to indicate a BIG problem, and you'll more than likely get the eventual advice that it has given up.

Too many hops

Another likely error message is one telling of "too many hops". This is not an indication of someone along the way consuming a strong ale. It means that the email has been transferred from one server to another and another and has done this too many times without reaching its destination.

A common cause is that someone has mail being forwarded to an address which already has mail being forwarded to the other address -- in other words, an endless loop. However, there can be cases caused by communication problems, such as if several mail transfer servers in a region go down at the same time, so it is difficult to find a route which works. In most cases the direct connections are between the backbone connections of the sender and recipient, so, while an email may seem to have many hops, you are likely to find that these have some business connection between the sender and recipient.

For example legitimate hops can be: sender to his or her ISP, to the domain mail exchanger, to the firm owning the computer where that domain is hosted, to the company providing the rackspace for that computer, to the backbone provider, and then back up a similar line to the recipient's computer. That may seem a lot, but effectively does not go to any strangers.

Open Relays are generally now forbidden by anyone who wants to stay in the internet business. If you want to send mail via a domain computer which is not part of your local ISP's setup you will have to prove who you are by providing identification. This is most commonly achieved by forcing you to download mail waiting for you before you are allowed to upload any messages. This is controlled by the "authentication" settings in your email program and can be as simple as telling it to check mail before sending.

BlahBlahBlah

If email is taking a long time to arrive it can be useful to look at the headers which show just how the message was routed. You will not normally see these but there will be a setting, often under options, which lets you select "all headers". In Outlook Express choose File > Select Properties > Details and in recent versions of Eudora there is an option button named "BlahBlahBlah".

You will have to make allowance for time zones but most will show time differences of just a few seconds or at most a few minutes. Where I have found problems they tend to be within an organisation rather than between mail nodes -- a major service provider taking hours or even days to transfer a message within its own network, sometimes between cities but sometimes too between mail servers which are possibly on racks within a few feet of each other.

There have been examples in Australia of mail taking days to arrive and the ISPs involved issue apologies and say it will not happen again. However, it does seem that large attachments can be a problem...and is one I mentioned in my article on sending large files in PCUpdate for April 2005. (Basic answer: there are better ways to send large attachments).

Then there can be addresses or ranges of IP addresses which are blocked.

Let's take an example of an actual event. A fellow web host had trouble contacting one of his clients by email. The message bounced with an automated comment that the client's IP address had been blocked for abuse. But it was not the address provided by the web host which had been blocked. The client was still sending mail via his local ISP so the IP address which had been blocked

was one belonging to the ISP, and it had been blocked by the backbone provider for that ISP.

So a major ISP was blocked on at least some of its addresses by its major supplier. The person sending the message was close enough to phone the would-be recipient but had difficulty persuading him that the problem was not with the sender's tiny business and that it was very likely that the recipient was not receiving other mail. It was big and bigger playing the parts of Dumb and Dumber, but tiny hosting supplier lost a client.

Another ISP also decided for a while to block any email which even mentioned web addresses they thought of as being inappropriate including mentions of a site which had been mildly critical of the ISP.

Keep it plain

Whitelists are intended to ensure that a recipient gets email they want, but mail rejected by system filters does not get to a local stage of being seen by the whitelist filter. Also, some whitelist entries get a minus score to count against the plus scores of spam: while a score of 2 or 3 can be caused by "faults" such as having a generic address like hotmail or yahoo, or using a lot of HTML code, or a lot of images with few words, the counteracting score of being on the whitelist was a measly -0.1 (one tenth of a point) in one setup I saw.

I know several business owners who filter all email with HTML code to a folder which may get a cursory glance before they hit the delete button. HTML code with lots of nice display and with animated icons may be great for email to your aunt or nephew, but for business email to which you want a reply, keep it to plain text. There is a site at <http://www.expita.com/nomime.html> which tells how to send plain text in just about every program capable of sending email.

There was a problem with Outlook 2002 (since fixed) by which, when an attempt was made to send messages in plain text, the message was sent with an attachment called winmail.dat that confused more than a few non-Outlook users.

Your email program may also let you ask for a receipt either on getting to a possibly correct mailbox, or on being read, or even both. The problem here is that it may produce an annoying dialog box asking whether a response should be sent (I always click 'No') or it may put an email in the outbox which confuses senders who know they have not written a message.

Some setups are not thought through to the extent they should be. I came across online advice at a large college which had a frequently-asked-question of "My emails to the college are bounced back - what can I do?". There were a couple of suggestions and then the comment: If you would like to notify us about your problem please send an email.

All who use email should have a secondary address. That may be Yahoo or GoogleMail or, if normally using a webhost or domain address, the email address you will have been given by your local ISP.

Check the sender

There are hopes for the future. For example SPF (Sender Policy Framework). This is an extension to the domain records which tell systems where mail should be coming from. Unfortunately it has not been widely adopted. Every system that receives mail would have to implement some kind of SPF look-up mechanism into the mail server and that just has not happened.

Microsoft has modified SPF and called it CallerID for email and then SenderID Framework but it has not been generally accepted. Yahoo also has a system called DomainKeys that utilizes a public key stored in the DNS records. Again it has little acceptance. Maybe one day...

----

Info on email:

Lots of general hints: <http://email.about.com/>

Gmail will never die: <http://www.google-watch.org/gmail.html>

SourceForge documentation on email problems:

[http://sourceforge.net/docman/display\\_doc.php?docid=6695&group\\_id=1#overview](http://sourceforge.net/docman/display_doc.php?docid=6695&group_id=1#overview)

Sender Policy Framework: <http://spf.pobox.com/>

\*Gordon Woolf, a long time MelbPC member, is a former publisher and an author of several books who recently moved in to web hosting to cater for small scale publishers.