



## OSW / appiChar Technical Guide for IT Staff

### **Section 8: The Internet**

This section explains the benefits of networking, introduces basic concepts which are useful when thinking about a network and describes different ways of connecting to and linking networks.

#### ***Contents:***

- Connecting to the Internet
  - Choosing a connection*
  - Security*
  - Troubleshooting connections*
- Web caching and filtering
- Email
  - Internet email*



## Connecting to the Internet

### Choosing a connection

To make sure you get the right connection, consider carefully what you want to do with it. At the bottom end of the scale is the simple dial-up connection using a modem. The options then work up through ISDN (Integrated Services Digital Network), DSL (Digital Subscriber Line) or Cable and then leased lines. As you'd expect, packages from different suppliers will provide varying levels of service so you'll need to make sure you can get the options you require.

These days there are very few reasons not to use DSL in preference to a modem connection if you're in an area where DSL is available. It provides a good level of bandwidth at a fixed price making Internet access mostly painless and easy to budget for.

If you need more bandwidth than a modem can supply and you can't get DSL or Internet access from a local cable company then ISDN might be the best option. There are packages available which can provide an unmetered connection which simplifies budgetting. Leased lines may be the only option if you need more bandwidth than ISDN can provide or a better level of support than is available with DSL services.

When it comes to choosing an ISP, treat it as far as possible just like choosing any other supplier of services to your organisation. Offers which look too good to be true probably are!

### DSL

Almost all suppliers of DSL will give you an ADSL (Asymmetric DSL). If the local exchange has been the subject of Local Loop Unbundling by easynet (or eventually others) then you may have other options including SDSL. Whether this is important to you will depend on your application. Most Internet usage is asymmetric anyway – much more data is sent than received so ADSL is appropriate.



<http://www.adslguide.org.uk/> provides a great range of information on ADSL including technical details and ISP comparisons.

Apart from *easynet* and any other unbundling ISPs as they come along, any other ISP will actually be selling you a service based around products from British Telecom who actually provide the bandwidth between you and the ISP's network. When you place an order with the ISP they place an order with BT to get your line activated for ADSL. This will often go according to plan but not always and as with the ADSL service as a whole there is no service level agreement for implementation or fixes. If your connection is critical to the operation of your organisation then you may wish to consider a backup connection based on dial-up or ISDN.

### ISDN

In some ways you can consider ISDN to be a faster version of modem-based dial-up. Because it's a digital service, instead of a modem you have a Terminal Adapter or a router. You still dial to make a connection, either to an ISP or to another site if you want to transfer data directly. ISDN is still used by some organisations as a way of providing connectivity between sites although many now opt for VPNs (Virtual Private Networks) over the Internet, mostly for cost reasons.

### Leased Line

A leased line provides a permanent connection between two points. They are available, from BT at least, in a range of bandwidths from 64kbps to 622Mbps depending on requirements and available budget! Because it's a private circuit you don't have the same security concerns as if your traffic were routed over the Internet. With this option service level guarantees become available should they be required.

### Required hardware

Whatever type of connection you have, you'll need some way of attaching it to the rest of your network. The simplest way will usually be to obtain a router with the correct interfaces. Assuming you don't opt for a managed service which includes one, routers for use with ADSL start from around £50 for a very basic model. Prices then go up with quality and features.

## Security

Some sites seem quite happy to be connected to the Internet with little or no security. This is not always because they have considered all the risks before taking that decision!

If you're using NAT this will be providing a certain amount of protection from external people trying to get in to your network – apart from any port forwarding you have set up it won't be immediately possible to attack the internal network.

A bigger risk comes from *trojan horse* programs which initiate a connection to an outside machine from within the network, allowing access from the outside because the initial connection comes from inside.

At the very least, some form of port filtering should be configured on any public facing (i.e. attached directly to the Internet) devices, blocking any traffic which is not known to be required. Normally everything should be blocked and then you might allow ports 25 (SMTP – mail transfer), 80 (web browsing), 110 (POP3 email) and 443 (secure Web browsing) plus any other specific ports you know you need.

A firewall is better in that it should be able to inspect the packets which pass through and make decisions based on the actual packets contents rather than just the ports being used.



For a good Q&A on Firewalls see  
<http://www.vicomsoft.com/knowledge/reference/firewalls1.html>

Just because you think you're not a target doesn't mean you can do without any form of security. Many attacks are actually automated, scanning a range of IP addresses and probing each address to see if there's anything interesting or at least potentially attackable there.



GFI produce a very useful utility which allows you to scan for potential vulnerabilities. It's not perfect but does generally give useful advice on what it finds.  
<http://www.gfi.com/lannetscan/index.htm>

## Troubleshooting connections

Firstly, make sure everything is plugged in properly and switched on! Apparently obvious, it's amazing how often this isn't done and thus how much time and hassle could be saved.



At this point you may wish you had a note of which lights should be off/flashing/on for each bit of equipment you have, as well as detailed configuration information. Remember to start on this as soon as you have fixed the current problem!

Remember to attack the problem methodically, where appropriate using the layered model developed earlier. Work out where traffic is going until you can establish the nature and location of the problem.

Start with making sure you can see other machines on the local network, including the router, using `ping`. Try to ping an address or two out on the Internet, both by name and just by IP address. You'll often have been given the IP address of your ISP's DNS servers so these can be a good start if you can't think of anything else. If ping by address works but not by name then you may "just" have a DNS problem – check which server you're using (with `IPCONFIG /ALL` for example) and perhaps change it to a different one if you have any alternatives.

If things are OK on the internal network but not externally, try using `tracert` to see how far packets are getting. If you get more than about three lines of response then packets are getting out of your network but are failing somewhere else. Try to `tracert` to a different address and see if the problem stays in the same place.



It is possible to configure some routers and firewalls so that `ping` and `tracert` won't work. If yours has been configured this way it won't be possible to use `ping` and `tracert` to try and resolve problems.

If packets don't appear to be getting anywhere you'll need to check that the whatever device you have connecting you to the Internet is on and connecting OK. Exactly how you do this will obviously depend on the device you have. Many will have a WAN or ADSL light or some other indication that the line they are connected to is active. It may be worth talking to your ISP at this point to see if they can shed any light on the issue – they should at the very least be able to tell you if they can see your connection from their side.

## Web Caching

When sharing an Internet connection between a number of users you may find there are a number of sites which many of them visit. In order to minimise the impact on your available bandwidth you can implement a cache which intercepts any web page requests and keeps a local copy. This means that the first person to visit a site will find it takes the “normal” length of time to download but subsequent users should find pages (which have been visited by other users) display more quickly than would otherwise have happened.

Not all content can be cached. Dynamic content, created when a user visit a dataases driven page or performs a search on a site perhaps, cannot usually be cached successfully. Also, secure http (URLs starting https) cannot be cached since it's obviously only appropriate for the one user at the time they're visiting.

Packages which provide Web caching include Microsoft ISA Server (on Windows) and Squid (on Unix)

## Web Filtering

There are many reasons for wanting to filter where you users can visit. Typically it's to stop them visiting content which is considered “inappropriate” in some way or other. The key problem with most filtering technologies is how to keep filtering that which should be filtered whilst still allowing access to that which shouldn't be. In some situations, “false positives” can be a big problem. For example, you might not intentionally want to block a site about breast cancer but sometimes it will be simply because the content includes breasts.

Since one of the categories of site which people often want to filter are to do with sex, sites dealing with sexual heath are also often affected. Because new sites are appearing all the time, keeping filters up to date can be a very time consuming and/or expensive business.

Packages which provide Web filtering include Microsoft ISA Server, SurfControl and many others (on Windows) or SquidGuard (on Unix, requires the use of Squid). There are also a number of network appliances which provide filtering.

## Email

There are a few ways of providing email to users and a number of mechanisms which are used to transport it, depending on the context. As far as getting mail to and from clients is concerned there are three main options:

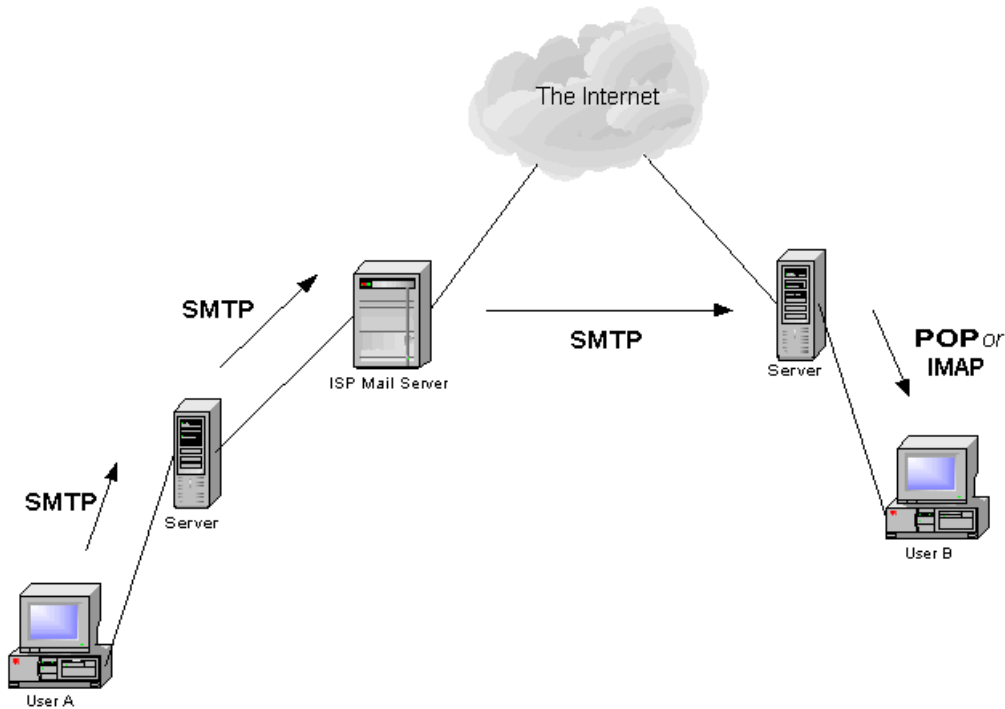
- SMTP (Simple Mail Transport Protocol) with POP (Post Office Protocol) or IMAP (Internet Message Access Protocol)
- Proprietary transport (e.g. as used by Microsoft Exchange)
- Web-based email

The presence of Exchange will normally dictate that Outlook will be used as the client, or possibly Outlook Web Access. Any problems which arise will need to be tackled using Event Viewer and enabling any diagnostic options that are available. You should also find help at the Microsoft web site for Exchange, currently at: <http://www.microsoft.com/exchange/support/default.asp>

Solving problems with a Web-based email system will depend on knowledge of the server and front-end. The usual troubleshooting rules apply – work methodically, isolate the problem and attempt to establish why it's happening. Or better yet, pass the problem over to someone else who has a greater knowledge of the specific system.

## Internet Email

Mail on the Internet will be transferred using one of SMTP, POP or IMAP depending on where the mail is going from and to. The diagram below shows how mail may be transferred between two users. Most of the process is normally quite transparent.



When User A's mail client send a message it uses SMTP to send it to a server on the local network. This server makes a decision based on whether the mail is local, in which case it is delivered to a mailbox on the same server or for anywhere else in which case it's sent (using SMTP again) to the ISP's mail server. The ISP's mail server will use the DNS MX record for the destination domain to decide which server to send the mail on to.

The advantage of using an ISP's mail server (a.k.a. *mail relay* or *smart host*) is that the ISP's server is "local" and you should be able to get a fast, reliable connection to it. It can then deal with all the looking up of DNS data and sending mail on to the final destination. The disadvantage is that you lose a small amount of control but this is usually only important when things are going wrong and most mail server software will allow you to select delivering directly instead of via a relay which can be useful for testing.

Once the message gets to the final destination server it will be delivered to some kind of mailbox until User B collects it using either POP3 or IMAP4. Note that unlike the previous steps where mail is pushed to the destination, for the final leg the recipient actually has to go and collect the message(s). Most clients will let you schedule this to happen automatically every so many minutes.

## Internet Connections and Applications

At each stage in the message sending process, the servers involved should add an additional line of information to the header of the message which is a part not normally displayed by email clients. For example, the headers for the message above would grow something like this:

### Initial Message

To: [userb@otherplace.org.uk](mailto:userb@otherplace.org.uk)

From: [usera@myplace.org.uk](mailto:usera@myplace.org.uk)

### When it gets to the server it becomes:

To: [userb@otherplace.org.uk](mailto:userb@otherplace.org.uk)

From: [usera@myplace.org.uk](mailto:usera@myplace.org.uk)

Received: from usera by server.myplace.org.uk with local (Exim 3.22 #1) id 193bIc-0006Ei-00; Thu, 10 Apr 2003 13:37:02 +0100

### And then on to the next step, at the ISP mail server

To: [userb@otherplace.org.uk](mailto:userb@otherplace.org.uk)

From: [usera@myplace.org.uk](mailto:usera@myplace.org.uk)

Received: from usera by server.myplace.org.uk with local (Exim 3.22 #1) id 193bIc-0006Ei-00; Thu, 10 Apr 2003 13:37:02 +0100

Received: from server.myplace.org.uk (server.myplace.org.uk [192.168.40.60]) by mail.megaisp.com (8.9.3p2/8.9.3) with ESMTTP id NAA14205 for <[userb@otherplace.org.uk](mailto:userb@otherplace.org.uk)>; Thu, 10 Apr 2003 13:37:05 +0100

### Until it finally reaches the mail server which holds the mailbox for User B.

To: [userb@otherplace.org.uk](mailto:userb@otherplace.org.uk)

From: [usera@myplace.org.uk](mailto:usera@myplace.org.uk)

Received: from usera by server.myplace.org.uk with local (Exim 3.22 #1) id 193bIc-0006Ei-00; Thu, 10 Apr 2003 13:37:02 +0100

Received: from server.myplace.org.uk (server.myplace.org.uk [192.168.40.60]) by mail.megaisp.com (8.9.3p2/8.9.3) with ESMTTP id NAA14205 for <[userb@otherplace.org.uk](mailto:userb@otherplace.org.uk)>; Thu, 10 Apr 2003 13:37:05 +0100

Received: from mail.megaisp.com (mail.megaisp.com [192.168.40.60]) by server.otherplace.org.uk (Postfix) with ESMTTP id D407734BD7 for <[userb@otherplace.org.uk](mailto:userb@otherplace.org.uk)>; Thu, 10 Apr 2003 13:37:30 +0100

An email client will normally have some facility which allows you to view the headers for a message, possibly through some sort of properties or view options menu. It's worth having a look at a few just to get a feel for how they work – there's much more information than shown above!