Demystifying IT

Information Professionals often develop very in-depth knowledge of specific areas of IT, which tend to be the areas they've had to battle with to deliver their Information Services. But a lack of understanding of how the whole of the IT infrastructure in an organisation works means that each IT issue they have to deal with is a new uphill struggle.

Let's start with your desktop computer.

For most people, their desktop computer is the only part of the IT infrastructure they really think about. But on a corporate network, the desktop is just on the periphery, and is relatively unimportant. Any decent IT network is designed so if your computer fails, your information doesn't get lost because it's all stored elsewhere.

All computers these days are **networked**, and it's this which makes the desktop so unimportant. At the most basic you connect two or three computers together with a network cable and a piece of equipment called a switch.

As things like printers, networked photocopiers, phones and so on became more intelligent, they tended to be connected to the network too. The de-facto standard for physical networking is called **ethernet**, which specifies the cables and the electronics which use them. In addition, there's **wifi**, which replaces the cable with a base station and wireless connection.

The **switch** connects everything in the network. Switches are just small boxes, and every computer is connected to the switch so when one wants to talk to another, the switch sends the information between them. A network in a medium sized company will probably have about one switch between twenty computers, and then you connect the switches together to create a bigger network.

These interconnected switches, computers and other devices, all form a Local Area Network, or LAN. LANs connect things within a single building, allowing everything to talk to the servers.

The most important thing on your local network are the **servers**. Servers are very like your desktop computer. They've got all the same bits, but more of them and they're more expensive. The server is the most critical part of your IT infrastructure, because it's where all the information is stored and processed.

Servers are designed to be always on. Good IT administrators will have bought servers with redundant parts, so if anything fails the server doesn't stop, and there will be a monitoring system which alerts the administrator at the first sign of trouble. Desk-side servers look like normal computers and live in normal offices. Rack mounted servers are long and flat and get slotted into big racks and live in data centres.

Everything on the network is orientated around the servers, with the relatively unimportant desktop machines asking the servers to do all the work. This principle is known as **Client/server architecture.** Your desktop is the client, where you view and edit the information you're working on.

The server stores and processes this information. So it's the server which is carefully tended to by your IT colleagues, because that's the vital part of the system. Client/server architecture used to be a revolutionary new thing with big debates about whether it was sensible, but nowadays, it's just the way everything is done.

You won't necessarily have very many servers in an organisation. It can run from none to hundreds for specialist applications. But a medium sized company might have three or four. Although a server is rather more of a serious computer than your desktop, what really distinguishes it is what software it runs.

Operating Systems and Software

So far we've just looked at hardware - the physical boxes in a network. But without software, these boxes are useless. We're going to look at the software which runs on desktops and servers. They're very similar, both

Copyright ONEIS Ltd 2015, All rights reserved. **www.oneis.co.uk**

with three layers of software.

Firmware starts everything going, and gets a machine running. It doesn't do all that much, and within a few seconds hands over to the operating system. For computers, firmware is often called the BIOS, "basic input/output system", and it's the bit which displays the white text on black which appears when you switch the computer on.

Once the firmware has got everything going, the **operating system** (OS) takes over. It's the bit which make the computer "operate" doing all the menial tasks like storing files and talking to the network.

On top of the OS, runs the **application** and **server software**, the reason for actually having a computer.

People often talk about **platforms**. This is a vague term: it could be the hardware, the operating system, or even a particular bit of application software.

It's most likely to refer to the operating system running on the server, and you'll hear names like Microsoft Windows, UNIX, Linux and Solaris. Each has different strengths, but choosing one is relatively simple -- it all comes down to what applications you want to run.

Sometimes you'll get a choice for a given application, in which case your IT people will choose the platform they're most familiar with.

So that's a brief run down of your Local Area Network - how all the IT equipment within your building talks to each other. Now let's look at how the computers within your building talks to computers outside your building.

When your company grows and opens a new office, the LANs in each building are normally connected together. They're connected together by a **WAN - a Wide Area Network**.

The key piece of equipment here is the **router**. These are a bit like switches, but they are more sophisticated. Instead of only being able to send information to computers or switches directly connected to them, they know how to reach any computer, by 'routing' information to the next router, until it reaches it's destination.

The line connecting the routers together will normally be leased lines, which are conceptually like very long but slow network cables you rent from a provider.

So that's connecting two offices together. But how about connecting to the rest of the world?

The **Internet** is an example of a WAN too. It's called the "Internet" because it interconnects networks. It started in the late 60s, and over the next few decades moved from being a largely academic network to the massive network which seems to have invaded every part of our lives.

The Internet is actually quite simple, and uses the same bits of equipment you find in your office. Again it is the routers that enable everything to talk to everything else.

The router in your office connects to an Internet Service Provider, or **ISP**. The ISPs then connect their routers together, so you can get to any other computer on the Internet, and all the servers in data centres.

In the domestic market, for home internet connections in the UK, ISPs include BT, Pipex, Virgin and Tiscali. But organisations chose more specialist ISPs, including BT, but there are lots of business-only ISPs. Which one you choose will depend on how reliable you want it to be, how many people will be using it, and how much you're prepared to pay.

Your internet connection arrives in one of two ways, a leased line, ADSL, newer DSL variants like VDSL, or even fibre. A leased line is the same thing you'd use to connect offices, but is rather expensive. So smaller offices and homes use DSL (Digital Subscriber Line.) It's a way of getting a fast-ish network connection over a standard phone line, and is a pragmatic way of reusing the existing telephony infrastructure. Most "fibre" connections sold today don't actually provide fibre to the premises, but terminate in a cabinet nearby and use DSL to deliver it to your home or office.

Copyright ONEIS Ltd 2015, All rights reserved. **www.oneis.co.uk**

Sending of information from your computer to another computer over the internet is not instantaneous. There's a physical limit on how fast packets of data can travel from one computer to another over a network connection. That limit is known as **Bandwidth:** which is a measure of how quickly information can be sent, measured in Mb/s. Each different type of network connection has a different bandwidth.

Your computer sends information to the router, which sends it on to the internet. This uses a leased line, or DSL link, which is much slower.

This is why your IT department discourage you from sending large emails and watching videos. The links to the internet are expensive and relatively low bandwidth, so if everyone uses them all at the same time, everything slows to a halt.

A key part of the Internet is **data centres**.

These are buildings full of servers offering public services you might access over the internet like email, websites, databases, and information services like Lexis Nexis. Some companies which need to run lots of servers may have their own private data centres.

They contain rows upon rows of servers, all installed in racks. They are basically a warehouse full of computers. They're useful for many reasons. Lots of ISPs have equipment there to inter-connect between themselves, so it's comparatively cheap to get a very fast internet connection. This is handy if you're running high traffic web sites. Data centres are built with security in mind and guarded 24 hours a day. And finally, they're the perfect place to run your servers, with plenty of air conditioning to keep them cool and backup power generators in case something goes wrong with the electricity supply.

Network addresses

You've got all these computers on the internet, and they need some way of specifying which other computer they want to talk to. Just as every house has a postal address or zip code used to deliver letters, computers have Internet Protocol addresses which routers use to deliver information.

You've probably come across IP addresses. This is what they look like, in a form called "dotted quad": 192.168.99.201

Every computer needs to have a unique address so that information can be routed between them on the Internet. However, there's only four billion of these IP addresses, which isn't enough because of the sheer number of devices connecting to the internet.

So organisations tend to have a few unique **public** IP addresses which have been assigned to them, by their ISP if they're small or a regional registry like RIPE if they're big. And then within their network they can use whatever **private** IP addresses they want for all their individual pieces of equipment. Private addresses are just like public addresses, except they're in one of three ranges set aside for them.

Now there's a problem with using private IP addresses for computers which want to talk to the internet. Because everyone uses the same set, if private addresses were allowed on the Internet, servers wouldn't be able to send back the information because the "return address" could refer to potentially thousands of computers.

So we use Network Address Translation, or NAT, which runs on your router. This translates the private IP addresses into the public IP addresses which work on the internet.

For example, your desktop computer could have address 10.0.0.12 on your LAN. The LAN has a router running NAT, which might have a public address of 17.0.0.29. This uses the public address to communicate on the desktop computer's behalf with the rest of the world, translating the address from private to public. When replies come back, it does it in reverse.

But this is only one way. The reverse translation is only possible because the NAT device knows which computer initiated the connection. The rest of the internet can't connect to the desktop computer, because it

Copyright ONEIS Ltd 2015, All rights reserved. **www.oneis.co.uk**

doesn't have a public address.

If this sounds like it's a bit complex and liable to cause problems, then you'd be right. These IP addresses are legacy IPv4 addresses, and the current internet protocol is IPv6. This has much larger addresses, so many that they're effectively infinite. Every computer gets it's own address, and NAT isn't required.

However, because NAT exists, and deploying IPv6 requires investment, ISPs and corporates have been reluctant to use it. However, now that all the IPv4 addresses have been assigned, there is a very slow movement to deploy IPv6.

Security

In a perfect world, no one would make mistakes, and everyone would act with the best of intentions. In the real world, mistakes are made and bad people exploit them. This means your IT team are constantly doing battle to keep your network safe.

There are two sets of people who make these mistakes. Firstly developers, who write software, and secondly system administrators, who sometimes set things up incorrectly.

All software has problems. There are millions upon millions of lines of code which make up your computer and the applications running on it. A mistake in any one of those could be exploitable.

Systems administrators can make mistakes too, accidently getting the wrong permissions, or letting their users set passwords which are easily guessed.

All this means you don't want just anyone connecting to your servers and desktop computers, because there's a possibility that mistakes can be exploited.

So you use **firewalls** to stop them. Your firewall lives on your router. Network Address Translation goes a long way to prevent problems, because you can't connect to anything behind it. This is good enough for the home and small offices, and when you're using IPv6, a simple firewall which blocks unexpected incoming connections works well and is enabled by default in your router.

With more complicated setups, you might have rules in place to prevent any connections other than those you know are allowed. You'd often also run a packet filter. This is a piece of software that analyses the information packets coming into the LAN and stops any that it thinks are malicious.

Firewalls, routers and NAT are often combined into a single box for a small network. When you build bigger networks, you might separate them out.

So let's put this altogether, and look at how it all works when you browse a web site:

- your web browser asks the operating system to make a connection to the web server
- the OS uses the network hardware to talk to the network
- which transmits the request to the server
- the server hardware receives the transmission
- which passes it to the operating system
- which decodes it and asks the web server to respond
- the web server sends the information back, in the same way it came.

Let's look at what happens in the network in more detail:

- your computer wants to request a web page from a web server, which lives in a data centre on the internet
- it sends the request to the switch
- which sends it on to the router
- which sends it to your ISP

Copyright ONEIS Ltd 2015, All rights reserved. **www.oneis.co.uk**

- your ISP works out where the web server is located, and sends the request on via the internet to the ISP which connects the server to the internet
- and that ISP sends the request to the web server
- it responds, sending the reply back, the same way it came.

No magic

I often find that people assume that they'll never understand computers, and this stops them from getting the most out of something which is fundamental to their work, and all for no good reason.

We've looked at the setup behind your desktop computer, that lets you browse the web, access hosted information providers' services, send and receive email.

It is a bit complicated, but it's not magic.

Computing is logical. Each part of the network is doing quite small tasks, and everything builds on something else. The complexity comes through three or four decades of building things on top of other things.

So when you have an IT problem, approach it logically. Remember it's not magic, and problems can be solved with rational thought. With the problems you face, you're really dealing with the simple bits which underlie it all, and your only challenge is to work out which bit is at fault.

So here are a few tips on being your IT department's best friend when reporting issues:

- describe the problem clearly and comprehensively, including any error messages you see
- explain what you did, what happened, and how that differed from what you expected to happen
- don't jump to diagnose: this can cause you to miss out critical parts of information
- send screenshots if appropriate
- state which browser you're using and which version if relevant
- state the urgency (honestly!)
- use your IT department's ticketing system if they have one and file each issue separately
- don't panic! It's always easiest to solve problems calmly.

Copyright ONEIS Ltd 2015, All rights reserved. **www.oneis.co.uk**