technical overview

a guide to the technologies and standards underlying the delivery of IP services and applications

What is IP?

CABLE & WIRELESS

IP technical overview 1

Contents

Details of where to get more information on the subjects covered in this book are given on the inside back cover.

1 The basics

What is IP?

IP stands for **Internet Protocol**. It is the most basic protocol used on the Internet, and defines a unit called a **datagram**, the format in which information is transmitted over the network. Any information that can be 'packetised' into IP datagrams can be carried over an IP network; this includes voice and video as well as data.

IP is a 'connectionless' service — the datagrams forming a message or other communication may take any route through the network between their origin and one or more destinations, where they are reassembled on arrival. Each datagram must therefore contain **addressing** information. Datagrams may also contain information about the type of service being transmitted and the **Quality of Service** (**QoS**) the service expects to receive from the network. For more information, see Section 3.

Why has IP become so important?

IP is not a new technology; it has been around for longer than frame relay and ATM. So why has it so dramatically taken off in importance?

In summary, the main driving factors are:

- the popularity of the Internet and browser technology
- businesses' need to stay competitive, by innovating and reaching new markets, and by rethinking the way they manage their networks

And the main technological enabling factors are:

- the development of very reliable, high-bandwidth public networks
- increasingly sophisticated terminal equipment, as processing power becomes cheaper and more plentiful
- a new generation of routers, designed for operation in large-scale, high-performance networks

Business drivers — a brief history

The widespread adoption of Internet technology

The Internet's growth in popularity was in part a result of the simplicity and ease of use of web browser software. Businesses also saw the advantage of having a single, consistent user interface for their internal 'enterprise' networks, and introduced Internet technology, creating **intranets**. The next step was to connect the intranet to the outside world — creating the **extranet**, with controlled access for external users, such as suppliers, customers and partners.

The Internet has now become a global marketplace, with the rapid advance of **e-commerce**, often referred to as B2B, for business-to-business, and B2C, business-to-consumer. E-commerce applications have been developed in parallel with technologies for secure transactions and more reliable end-to-end communications.

The need for flexibility

Organisations need their networks to be flexible, so that sites can be added and removed easily and at short notice. They also want any network site or device to be able to communicate directly with any other one, in a **meshed** network.

Services like frame relay and ATM, with permanent virtual circuits (PVCs) between sites, sharing access bandwidth using statistical multiplexing, were an improvement on the 'traditional' solution based on private circuits. But with these services, adding a new site still involves reconfiguration throughout the network, and access bandwidth is a limiting factor. With IP, however, a new site or device can be added to the network very easily — no reconfiguration is needed and access bandwidth availability does not constrain the number of connections.

The trend towards virtual private networking

As companies' enterprise networks become more inclusive and wide-reaching, the challenge of day-to-day management can be a distraction from the strategic planning and innovation necessary for ongoing competitiveness.

Businesses are increasingly turning away from fully 'private' networking solutions, and relying on operators like Cable & Wireless to supply and manage their infrastructure, allowing them to:

- reduce capital investment and management overheads
- concentrate their own resources on developing and supporting applications
- benefit from the performance, reliability and capacity of a public network, and the expertise and round-the-clock availability of specialist staff
- be confident that future upgrade paths are built into the solution, mitigating technical risks

This is the **virtual private network (VPN)** — a solution providing secure transport of private traffic across public network infrastructure, with **Service Level Agreements (SLAs)** defining availability and performance. IP-compatible VPNs are known as IP-VPNs.

Technological enablers

Advances in public network quality and capacity

Advanced public networks, based on optical fibre and synchronous digital hierarchy (SDH) transmission, are now very reliable and capable of carrying huge amounts of data. The use of dense wavelength division multiplexing (DWDM) is greatly increasing the accessible bandwidth of the fibre infrastructure, and the introduction of optical switching technologies will further increase network capacity.

Cable & Wireless is at the leading edge of network development, investing more than $3.5 billion worldwide in the infrastructure to support seamless IP solutions for business — a managed global SDH network, with 99.99%+ availability and fully scalable architecture.

When networks were less reliable, they needed to incorporate error detection and correction at intermediate points on transmission paths, to ensure end-to-end integrity. Most of this 'network intelligence' can now be moved out from the network core to the edges — powerful terminal devices at access nodes or on customers' sites. In terms of the OSI seven-layer model, the network can operate at a lower level of the protocol stack (see the Appendix and the diagram on page 18).

Advantages of this configuration include:

- Because the edge devices are developed independently of the core, new services and applications can readily be added and new features are easily tested.
- It is relatively easy to increase the capacity of the core, simply by adding more bandwidth as traffic levels demand.
- The capabilities of connectionless technology like IP — for example, its use for **broadcasting** information to multiple destinations — can be fully exploited, with plentiful bandwidth, and intelligent equipment taking care of security and data integrity.

New-generation routers

Although more flexible than switches, routers were until recently seen as potential bandwidth bottlenecks. However, the new generation of routers has been designed for use in large-scale, high-performance IP networks. Based on modular, scalable architecture, these routers can grow in capacity by increments. They support very fast transmission and interface rates — **terabit routers**, with aggregate speeds of 10^{12} bit/s and above, are now becoming available.

2 Key benefits of IP

Ubiquity

IP is everywhere — on users' desktops, at their workplaces, in their homes and on the move. Browser software is easy to use and gives a single, consistent interface to all kinds of information with no need for special training. The range of access devices is growing, notably with the introduction of mobile phones using **Wireless Access Protocol (WAP)**.

Anyone with a connection to the public network — i.e. anyone with a telephone line, fixed or mobile — can be given access to an IP network. Sites which already have frame relay or ATM connections can use them for access.

Unlike a telephone number, an IP address is location-independent and device-independent — users can access their mailboxes wherever they are, provided they have configured their equipment.

Scalability and flexibility

As explained on page 4, the architecture of public IP networks makes it relatively easy to add bandwidth as demand increases, without major reconfiguration.

Users, sites and devices can be added to or removed from an IP-based enterprise network at any time, simply by validating their user IDs and passwords. An IP network is, by definition, meshed — a single, straightforward connection to a public network Point of Presence gives any-to-any connectivity, with no need to worry about configuring virtual circuits to share access bandwidth.

And IP networks can be made as secure as the fully private alternative, with the added advantage of easier and more flexible control over which users have access to what information.

Potential for innovation

IP is a universal protocol, capable of handling all types of traffic — voice, data and video — over a common infrastructure. It has a simple, standardised applications programming interface (API), designed for cross-platform compatibility, to facilitate development of new applications.

The 'decoupling' of service provision from network architecture gives PTOs such as Cable & Wireless considerable scope for introducing innovative managed solutions and applications targeted at businesses' special needs.

Ending the dichotomy between voice and data also opens up the possibility for computer–telephony integration (CTI) applications, not only in call centres but on every user's desktop.

Economies of scale

To date, most organisations have treated voice and data networks as separate entities. This has worked well, but has incurred two sets of management overheads, and necessitated two delivery mechanisms to users' desktops.

IP offers the potential to run all applications over a shared core infrastructure, with associated economies of scale which PTOs will be able to pass on to customers. Within the building, a single cabling system and integrated terminals will deliver multimedia communications to the desktop.

Meanwhile, the IP network will continue to support a combination of access configurations, including legacy PSTN, frame relay and ATM, preserving customers' investment in equipment and circuits.

3 Technical issues

IP in the network

IP is a network-layer protocol — it operates at layer 3 of the OSI model. It is tightly integrated with **Internet Control Message Protocol** (**ICMP**), which provides diagnostic functions, generating messages to alert hosts of any unusual conditions within the network (a host is an 'intelligent' device connected to the network, i.e. one which is capable of transmitting or receiving data).

IP version 4 (**IPv4**), which has been around for many years, is still the predominant version in use. **IPv6**, also known as **IPng** ('next generation'), was introduced in 1998 and features:

- 128-bit addressing, using prefixes to simplify routing (see 'IP addressing' below)
- autoconfiguration ('plug and play' addition of new devices)
- new security features
- support for real-time communications and multicasting

IPv6 is designed to allow a gradual transition from IPv4, with the two versions working side by side. However, its adoption so far has been slow, partly because of its potential impact on the global routing tables that track IP addresses, and because older routers in the network will not recognise the address format.

IP and TCP

As a connectionless protocol, IP can be compared to the postal system — it lets you address a package and drop it into the system (with the equivalent of a first- or second-class stamp for QoS) but there is no direct contact between you and the recipient. Most networks combine IP with **Transmission Control Protocol** (**TCP**), a layer 4 (transport layer) protocol that allows hosts to establish a virtual connection and exchange data. TCP has some error-checking capability — when the datagrams making up a message are received and reassembled, the receiving host checks to see whether any parts are missing and asks the transmitting host to resend them if necessary.

The TCP/IP protocol suite became the *de facto* standard for LANs, being more practicable to implement than equivalent OSI protocols. However, before the Internet provided the impetus, TCP/IP was not extended into the wide area because of two perceived deficiencies — it could not give application-level QoS guarantees on this scale, and it did not offer adequate security assurances.

Voice and video in the IP network

The quality problem is further complicated when IP networks are used for real-time voice and video and for data streaming. With streaming, the information starts being heard and/or displayed at the receiving end before the whole file has been transmitted; so, if anything gets lost in transmission, there is no way to recover it.

Instead of TCP, real-time applications use a 'cut down' transmission protocol called **User Datagram Protocol** (**UDP**), which provides a direct method of sending and receiving datagrams over an IP network, without the error recovery overhead. Real-time applications therefore need the network itself to give quality assurances.

The challenge of IP network provision

With the drive towards IP networking, the challenge is to deliver guaranteed QoS and security for real-time and non-real-time services on a connectionless network without limiting flexibility.

Various standards and protocols exist and are being developed. For large-scale public IP networks, a technology called **Multi-protocol Label Switching** (**MPLS**), offers the greatest potential. MPLS is used in the Cable & Wireless IP network to provide customers with a high-quality IP-VPN service.

MPLS is discussed in more detail on pages 12–16.

IP addressing

Each device on an IP network has an identifier, or address, consisting of four eight-bit binary numbers (bytes), 32 bits in all. An IP address is normally written as four decimal numbers, each with a value of 0 to 255, separated by full stops, e.g. 123.156.78.212.

These four numbers are used in different ways to identify a particular network and the host on that network. Three classes of Internet addresses were originally specified, in the document RFC 791:

- Class A addresses use the first byte to identify the network and the other three for the hosts. Class A supports 16 million hosts on each of 126 networks. Class A addresses have all been assigned, mostly to large organisations who use **subnetting**.
- Class B uses the first two bytes to identify the network, and supports 65,000 hosts on each of 16,000 networks.
- Class C uses the first three bytes to identify the network, and supports 254 hosts on each of 2 million networks.

Partly because of past over-assignment of Class A addresses, the number of available IP addresses is running out, and the class-based system is gradually being superseded by **classless interdomain routing** (**CIDR**), a system which is tied to the adoption of IPv6. With CIDR, a single address can be used to identify many unique IP addresses, allowing **route aggregation** for more efficient routing.

Domain names provide a more easily remembered alias for IP addresses. Translation between the two is controlled by a distributed database system called the **Domain Name Service** (**DNS**).

Quality of Service (QoS)

QoS is a set of measurable parameters defining the performance of a network and forming part of the contract between the network user and the network provider. QoS measurements cover transmission quality, service availability, maximum permitted latency, etc.

Different applications have different QoS requirements — in particular, real-time traffic such as voice and video cannot tolerate transmission delays or packet loss. So ways must be found to give different traffic types differential treatment when congestion occurs; and to measure QoS performance in order to comply with customers' SLAs. One way to avoid congestion is simply to add bandwidth; this may not be a problem in the LAN, but on a larger scale additional mechanisms are needed.

QoS in the LAN and at the network edge

The edge of the network — access circuits — is typically where bandwidth is most limited, and where it is important to assign priority to real-time traffic. Two mechanisms used at this level are:

- **Real-time Transport Protocol (RTP)**, which allows network gateways to give some IP packets a special weight by modifying bits in the IP header.
- **Resource Reservation Set-up Protocol (RSVP)**, which can provide quantifiable latency and bandwidth, especially for multi-cast applications and high-bandwidth traffic flows. However, RSVP is not seen as practicable for voice traffic on a large scale.

DiffServe — QoS in the core

The DiffServe (differential QoS) working group of the Internet Engineering Task Force (IETF) is defining standards for marking IP data packets with traffic class and traffic contract information, so that they can be treated accordingly by the network. The type of service (TOS) field in the IP packet header has been redefined as the

DiffServe field, and information in it determines the **per-hop behaviour (PHB)** — the forwarding treatment a packet expects to receive at each network node.

DiffServe will support dynamic QoS, based on time of day, application or user identity.

Security

Security is crucial when confidential information is being transmitted over public network infrastructure, and secure transactions are essential to e-commerce. The **IPSec** working group of the IETF is defining standards to provide cryptographic security services with flexible support for:

- confidentiality (encryption)
- authentication (proof of sender)
- integrity (detection of data tampering)
- access control
- replay protection (preventing unauthorised resending of data)

The IPSec working group is also developing techniques for **Internet Key Exchange (IKE)**, specified in the **Internet Key Management Protocol (IKMP)**. IKE, which is based on public key algorithms, establishes keys for encrypting and decrypting information.

With IPSec, the entire contents of the IP packet are encrypted for security, and digitally signed for data integrity and authentication. It works in two ways:

Transport mode is the direct relaying of protected data between hosts with their own IPSec capability — for example, PCs with inbuilt client encryption.

In **tunnel mode**, the IP traffic is encrypted and authenticated by IPSec gateways at the entrance and exit points of the network. The sending gateway encapsulates the entire IP packet, including the

original IP header, with IPSec encryption. It then adds an authentication header (AH) before sending the packet to the recipient gateway, where it is decrypted.

MPLS — the key to large-scale IP networking

MPLS is seen as the key to widespread implementation of IP-VPNs, and to the full integration of IP with connection-oriented services such as ATM and frame relay. Cable & Wireless uses MPLS in the network core to provide customers with the secure, quality-assured **IP-VPN QoS** service.

MPLS effectively makes routers behave rather like switches, but with much more flexibility. It integrates layer 2 (data link) information about network links (bandwidth, latency, utilisation, etc.) with layer 3 routing (see the Appendix for more information about the layers in the OSI model).

How routers traditionally work

In a router-based network, IP datagrams (packets) move from source to destination by being sent from one router to another, in a series of **hops**. Routers co-operate with each other, using protocols such as Router Information Protocol (RIP) and Open Shortest Path First (OSPF), to construct **routing tables**.

Each time a router receives a packet, it uses the destination IP address in the header as an index, to look up the identity of the 'next hop' router in the routing table. This look-up process occurs at each hop on the path from source to destination.

The construction of the routing tables and their use for look-ups are, in theory, separate logical operations; but, in practice, they are very tightly coupled in routers. Figure 1 illustrates how these functions interact.

Figure 1. Router functions

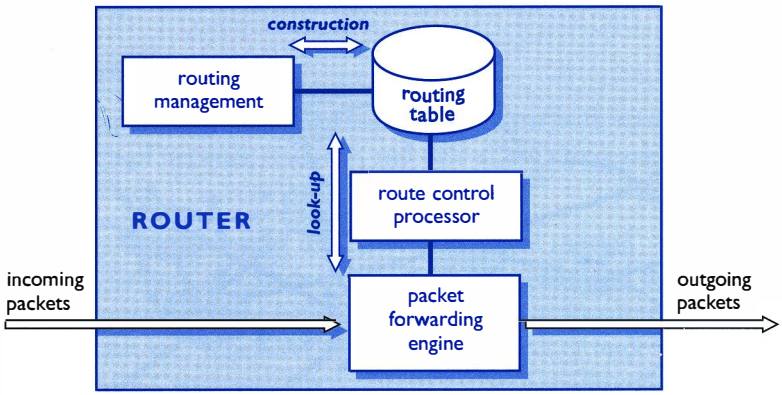

How MPLS works

The key feature of MPLS is the generation of short, fixed-length labels that act as a shorthand representation of IP headers. When a packet enters the network, a **label edge router** (**LER**) gives it a 32-bit label, based not only on the source and destination addresses, but also on the type of service and its QoS requirements.

According to their labels, packets are then assigned to **label-switched paths** (**LSPs**). High-speed **label switch routers** (**LSRs**) in the core network use signalling information to determine the routing of LSPs, taking into account network conditions (e.g. bandwidth availability). In essence, MPLS routers make *switching* decisions based on the labels. MPLS is an example of 'cut-through' processing, so called because packet processing overheads are reduced by associating multiple packets with a single data flow.

The IP header isn't used again until the packet leaves the network, when another LER removes the label.

The basic architecture of an MPLS-based network is illustrated in Figure 2 (page 14), which shows a fully scalable IP-VPN.

Figure 2. Customer sites linked in an IP-VPN using MPLS

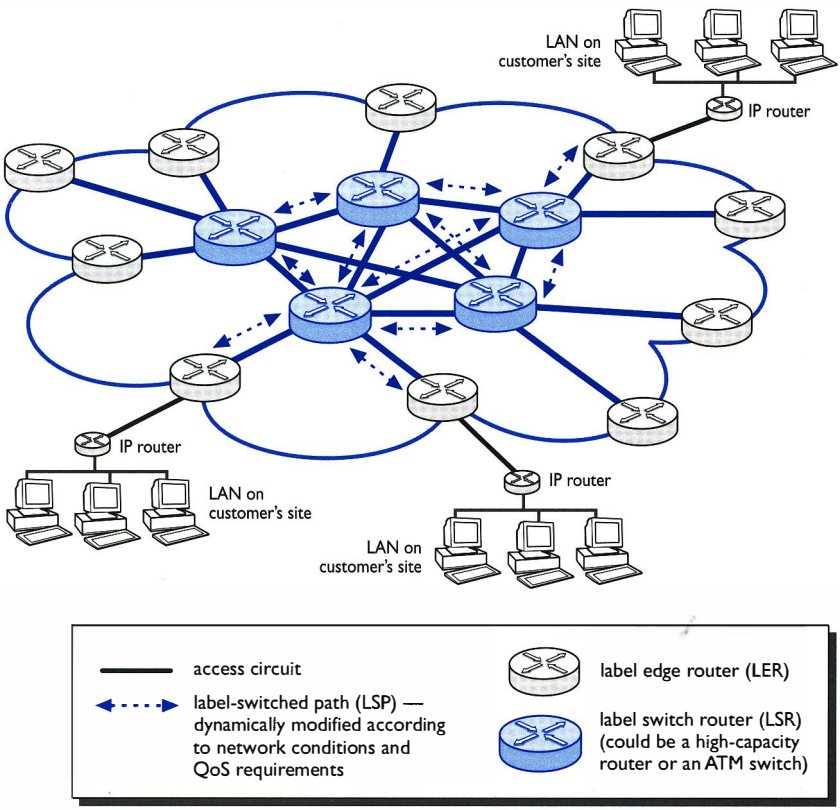

Why 'multi-protocol'?

The 'multi-protocol' part of the name comes from the facts that:

- Labels can be based on information other than IP headers, so MPLS can support other LAN protocols such as IPX.
- LSPs can be switched by ATM switches as well as by IP routers, allowing IP traffic to be carried across an ATM core with much less complexity than would otherwise be involved. This is

important if IP networks are to be successfully integrated with existing ATM networks.

Traffic engineering

Traffic engineering is the dynamic routing of traffic flows to optimise utilisation of network resources and deliver the QoS each traffic flow requires.

MPLS traffic engineering uses 'constraint-based routing' — the path for a traffic flow, or LSP, is the shortest path that meets that traffic flow's requirements for bandwidth, latency, priority compared to other flows, etc. Whenever congestion, or a link or node failure occurs, MPLS automatically adapts to the new conditions and modifies the path accordingly without disrupting end-to-end traffic flow.

Advantages of MPLS

MPLS gives network operators the flexibility to route traffic around congestion and bottlenecks, and to offer differential QoS to their customers. For example, customers who transmit or receive a large quantity of streaming media or large files can specify minimal latency and packet loss in their SLAs. At the same time, network complexity and costs are reduced.

Customers benefit from network performance equivalent to high-capacity private circuits, with higher resilience, flexibility and efficiency, and much lower costs.

MPLS can also help when the IP addresses used by organisations within their corporate networks are not globally unique. By encapsulating a non-unique address within a label which is unique as far as the MPLS devices are concerned, the need for complex address translation is reduced.

With the label-based forwarding plane (set of operations) effectively separated from the routing protocol control plane, the two can be independently modified. For example, there's no need to change the forwarding machinery when introducing a new routing strategy or

service into the network. Network operators therefore benefit from the potential for developing more diverse, reliable, competitive and innovative services.

Delivering Voice over IP (VoIP)

As the IP network is developed, voice and fax services will also be able to benefit from the capacity and economies of scale of a shared infrastructure. Integration in the core should be transparent to existing customers, so that they do not need to modify their telephone or fax equipment, nor to change numbers or access codes.

This is achieved by **media gateways**, which allow transmission of voice traffic across IP infrastructure by providing conversion between the information carried on telephone circuits and the data packets carried over IP networks.

Access services supported include:

- Basic Rate and Primary Rate ISDN
- analogue exchange lines
- mobile telephony
- H.323 access devices (H.323 is part of the ITU-T standards for audio, video and data communications across IP-based networks)

The call control intelligence for a number of media gateways is handled by a **media gateway controller** (**MGC**).

Communication between gateways and MGCs uses the **Media Gateway Control Protocol** (**MGCP**), defined by the Megaco working group of the IETF.

MGCs in the Cable & Wireless IP network also communicate, via SS7 signalling gateways, with the C&W Intelligent Network (IN), for provision of enhanced services to call centre operators, Internet Service Providers (ISPs) and other organisations offering IP-based facilities to customers, staff and associates.

Appendix: The OSI model

Open Systems Interconnection (OSI) is an internationally agreed model, developed by the International Standards Organisation (ISO) to enable interworking between systems worldwide.

In practice, actual OSI standards were never widely adopted. However, the OSI *model* is useful for understanding the processes involved in data transmission and for putting other internetworking standards and technologies in context. The original model is structured in seven layers, as shown in Figure 3, although today's networks effectively use a simpler protocol stack, with network infrastructure and end systems possibly including some or all layers.

The basic principles of the OSI model are:

- When data is transferred between systems, various processes must be carried out. These can be defined in terms of the 'logical distance' between a user's desktop or a piece of application software and the physical medium that carries the data.
- By arranging these processes in layers, we can define independent end-to-end protocols — called **peer protocols**.
- A key question is: 'At what point in the end-to-end transmission path are these processes taken care of?'

If the physical network is very reliable and conditions are predictable, it may be safe to leave tasks like error detection and correction to the equipment at each end.

A network operating at a very low level — clear-channel bandwidth — can carry data from all sorts of sources, independently of higher-level protocols. When higher-level network and transport protocols are used on the transmitting network, data must be encoded in some form for transmission — for example, into frames, ATM cells or IP datagrams.

Figure 3. The OSI model

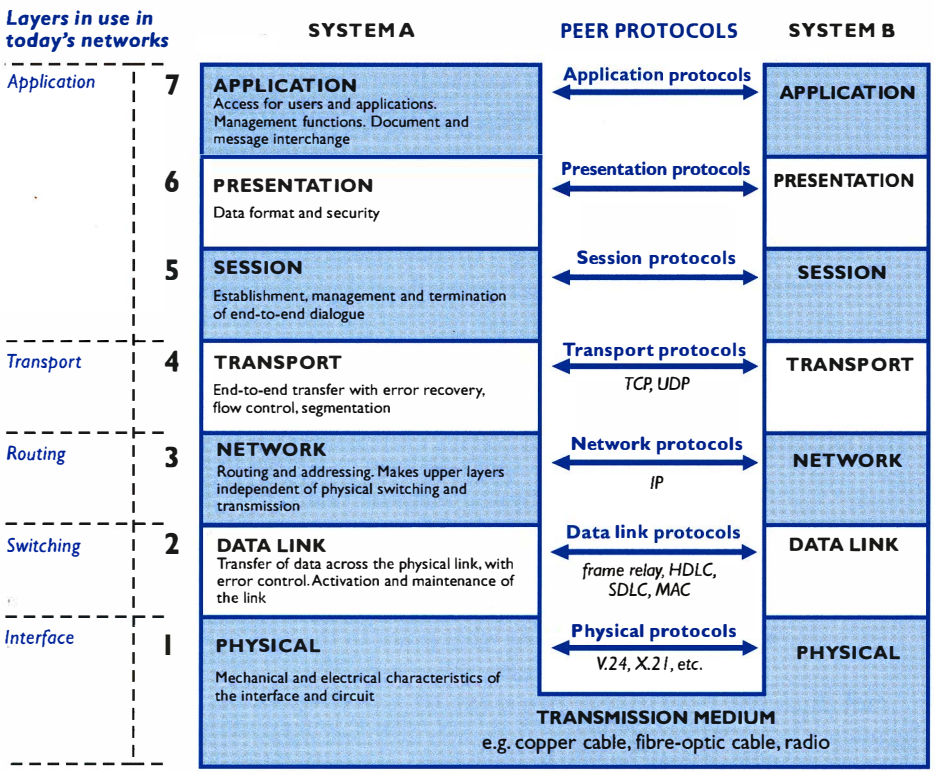

Switching and routing at layers 2 and 3

Figure 4 shows the positioning of different network nodes in the layered protocol hierarchy. The basic principles here are:

- A lower layer should not need to know about layers above it — e.g. a layer 2 switch isn't aware of information in IP headers.
- Nodes at higher layers have more information available to them than lower ones; so they can exert more control over forwarding decisions.

Switches traditionally operate at layer 2, the **data link** layer. With little processing overhead, they can forward traffic between network segments with high throughput and low latency, using 6-bit media access control (MAC) addresses. However, layer 2 switched networks have limited scalability and flexibility; they cannot always adapt to increasing complexity; and they have limited support for multi-media and multicasting.

Routers (see also page 13) operate at layer 3, the **network** layer, and make forwarding decisions based on information in individual packet headers. As well as address details, this information can include security features, such as access control, and differential QoS requirements. So routers can provide flexible, value-added services; being more 'intelligent', they are also better equipped than switches to handle increasing network complexity. But their processing overhead is higher, risking congestion if they cannot output packets as fast as they receive them ('wirespeed' forwarding).

Layer 3 switching aims to combine the speed of layer 2 switching with the intelligence of layer 3 routing. What we call layer 3 switches are, in effect, a new generation of *routers*, in that they may process each packet. But while routers carry out the processing in microprocessor-based 'engines', layer 3 switches use **application-specific integrated circuit** (**ASIC**) hardware, and can achieve much higher forwarding rates than conventional routers.

Another way to increase network speed while providing enhanced features and services is to use some form of 'cut-through' processing like MPLS (see page 13 for more information).

Figure 4. Inter-network nodes

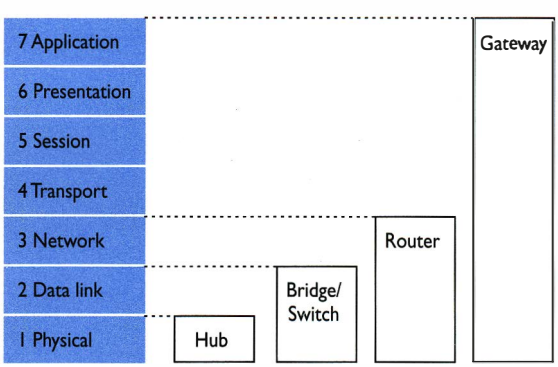

Index

More information

For more information on business solutions and services, speak to your Cable & Wireless sales contact, or visit the Cable & Wireless web site:

http://www.cwcom.co.uk

Descriptions of Cable & Wireless services such as IP-VPN QoS and Internet VPN can be found on the C&W consultants' web site:

http://www.cwcom-reference.co.uk/consult/index.htm

This includes full Service Descriptions, briefer Service Summaries and a technical Glossary covering many terms associated with IP networking and other communications technologies.

The companion to this booklet, ***IP applications***, gives an overview of web hosting and co-location facilities and features, and summarises some of the advantages of using Cable & Wireless as an IP applications supplier.

Registered office address:
124 Theobalds Road, London WC1X 8RX

JDS/MTP/BMK2075/07/00/V1