Internetworking

Note. The content of this document is mainly drawn from Wikipedia [www.wikipedia.org] and follows *GNU Free Documentation License (GFDL)*, the license through which Wikipedia's articles are made available.

The GNU Free Documentation License (GFDL) permits the redistribution, creation of derivative works, and commercial use of content provided its authors are attributed and this content remains available under the GFDL. Material on Wikipedia (and this document too) may thus be distributed multilingually to, or incorporated from, resources which also use this license.

Table of contents

1 INTRODUCTION

The *Internet*, or simply the *Net*, is the publicly accessible worldwide system of interconnected computer networks that transmit data using a standardized Internet Protocol (IP). It is made up of thousands of smaller commercial, academic, domestic, and government *networks*. It carries various information and services, such as electronic mail, online chat, and the interlinked Web pages and other documents of the World Wide Web.

Contrary to some common usage, the Internet and the World Wide Web are not synonymous: the Internet is a collection of interconnected computer networks, linked by copper wires, fiber-optic cables, etc.; the Web is a collection of interconnected documents, linked by hyperlinks and URLs, and is accessible using the Internet.

2 THE CREATION OF THE INTERNET

Lawrence Roberts was recruited to head a project to implement a network, and Roberts based the technology on the work of Paul Baran who had written an exhaustive study for the U.S. Air Force that recommended packet switching to make a network highly robust and survivable. After much work, the first node went live at UCLA on October 29, 1969 on what would be called the ARPANET, the "eve" network of today's Internet.

The first TCP/IP wide area network was operational by January 1, 1983 (this is technically *the birth of the Internet*), when the United States' National Science Foundation (NSF) constructed a university network backbone that would later become the NSFNet. It was then followed by the opening of the network to commercial interests in 1995.

Important separate networks that offered gateways into, then later merged into the Internet, include Usenet, Bitnet and the various commercial and educational X.25 networks such as Compuserve and JANET. The ability of TCP/IP to work over these pre-existing communication networks allowed for a great ease of growth. Use of Internet as a phrase to describe a single global TCP/IP network originated around this time.

The network gained a public face in the 1990s. In August 1991 CERN in Switzerland publicized the new World Wide Web project, two years after Tim Berners-Lee had begun creating HTML, HTTP and the first few web pages at CERN in Switzerland. In 1993 the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign released the Mosaic web browser version 1.0, and by late 1994 there was growing public interest in the previously academic/technical Internet. By 1996 the word "Internet" was common public currency, but it referred almost entirely to the World Wide Web.

Meanwhile, over the course of the decade, the Internet successfully accommodated the majority of previously existing public computer networks (although some networks such as FidoNet have remained separate). This growth is often attributed to the lack of central administration, which allows organic growth of the network, as well as the non-proprietary open nature of the Internet protocols, which encourages vendor interoperability and prevents any one company from exerting too much control over the network.

3 Today's Internet

Aside from the complex physical connections that make up its infrastructure, the Internet is held together by bi- or multi-lateral commercial contracts and by technical specifications or protocols that describe how to exchange data over the network.

As of January 2006, over 1 billion people use the Internet according to Internet World Stats (Fig. 1).

WORLD INTERNET USAGE AND POPULATION STATISTICS

World Regions	Population (2006 Est.)	Population % of World	Internet Usage, Latest Data	% Population (Penetration)	Usage % of World	Usage Growth 2000-2005
Africa	915,210,928	14.1 %	22,737,500	2.5 %	2.2 %	403.7 %
Asia	3,667,774,066	56.4 %	364,270,713	9.9 %	35.7 %	218.7 %
Europe	807,289,020	12.4 %	290,121,957	35.9 %	28.5 %	176.1 %
Middle East	190,084,161	2.9 %	18,203,500	9.6 %	1.8 %	454.2 %
North America	331,473,276	5.1 %	225,801,428	68.1 %	22.2 %	108.9 %
Latin America/Caribbean	553,908,632	8.5 %	79,033,597	14.3 %	7.8 %	337.4 %
Oceania / Australia	33,956,977	0.5 %	17,690,762	52.9 %	1.8 %	132.2 %
WORLD TOTAL	6,499,697,060	100.0 %	1,018,057,389	15.7 %	100.0 %	182.0 %

NOTES: (1) Internet Usage and World Population Statistics were updated for December 31, 2005. (2) CLICK on each world region for detailed regional information. (3) Demographic (Population) numbers are based on data contained in the world-gazetteer website. (4) Internet usage information comes from data published by Nielsen//NetRatings, by the International Telecommunications Union, by local NICs, and other other reliable sources. (5) For definitions, disclaimer, and navigation help, see the Site Surfing Guide. (6) Information from this site may be cited, giving due credit and establishing an active link back to www.internetworldstats.com.

Fig. 1 - INTERNET USAGE STATISTICS - The Big Picture

4 Internet protocols suite

Unlike older communications systems, the Internet protocol suite was deliberately designed to be independent of the underlying physical medium. Any communications network, wired or wireless, that can carry two-way digital data can carry Internet traffic. Thus, Internet packets flow through wired networks like copper wire, coaxial cable, and fibre optic; and through wireless networks like Wi-Fi. Together, all these networks, sharing the same high-level protocols, form the Internet.

The Internet protocols originate from discussions within the Internet Engineering Task Force (IETF) and its working groups, which are open to public participation and review. These committees produce documents that are known as Request for Comments documents (RFCs). Some RFCs are raised to the status of Internet Standard by the IETF process.

Some of the most used application protocols in the Internet protocol suite are IP, TCP, UDP, DNS, PPP, SLIP, ICMP, POP3, IMAP, SMTP, HTTP, HTTPS, SSH, Telnet, FTP, LDAP, SSL, and TLS.

Some of the popular services on the Internet that make use of these protocols are e-mail, Usenet newsgroups, file sharing, Instant Messenger, the World Wide Web, Gopher, session access, WAIS, finger, IRC, MUDs, and MUSHs. Of these, e-mail and the World Wide Web are clearly the most used, and many other services are built upon them, such as mailing lists and blogs. The Internet makes it possible to provide real-time services such as Internet radio

and webcasts that can be accessed from anywhere in the world.

Some other popular services of the Internet were not created this way, but were originally based on proprietary systems. These include IRC, ICQ, AIM, and Gnutella, although all of those mentioned now have Free implementations, which in some cases are the most commonly used.

The **internet protocol suite** is the set of communications protocols that implement the protocol stack on which the Internet and most commercial networks run. It is sometimes called the TCP/IP protocol suite, after the two most important protocols in it: the *Transmission Control Protocol (TCP)* and the *Internet Protocol (IP)*, which were also the first two defined.

The internet protocol suite — like many protocol suites — can be viewed as a set of *layers*, each layer solves a set of problems involving the transmission of data, and provides a well-defined service to the upper layer protocols based on using services from some lower layers. Upper layers are logically closer to the user and deal with more abstract data, relying on lower layer protocols to translate data into forms that can eventually be physically transmitted.

The *OSI model* describes a fixed set of seven layers that some vendors prefer and that can be roughly compared to the IP suite. This comparison can cause confusion or give further insight into the IP suite.

4.1 Layers in the internet protocol suite stack

The IP suite uses encapsulation to provide abstraction of protocols and services to different layers in the stack. The stack consists of four layers:

4. Application

3. Transport

2. Network

1. Link

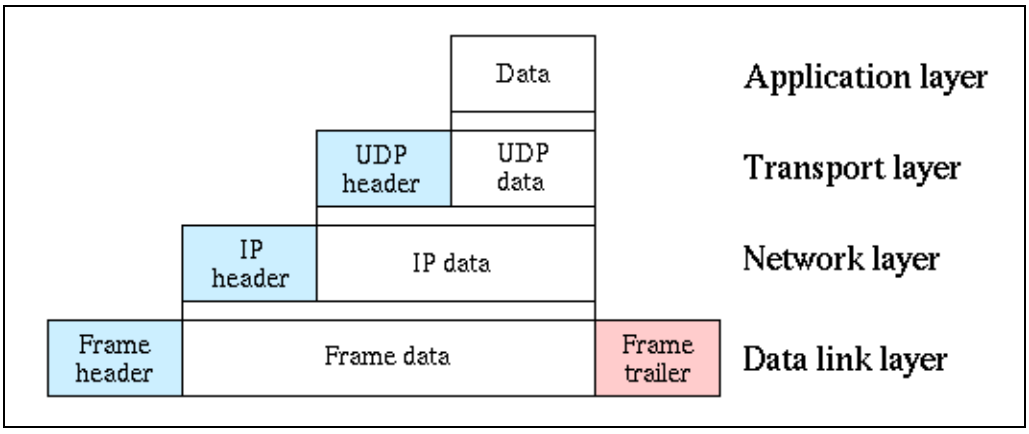

Fig. 2 - Sample encapsulation of data within a UDP datagram within an IP packet.

The layers near the top are logically closer to the user while those near the bottom are logically closer to the physical transmission of the data. Each layer has an upper layer protocol and a lower layer protocol (except the top/bottom protocols, of course) that either use said layer's service or provide a service, respectively. Viewing layers as providing or

consuming a service is a method of abstraction to isolate upper layer protocols from the nitty gritty detail of transmitting bits over, say, ethernet and collision detection while the lower layers avoid having to know the details of each and every application and its protocol.

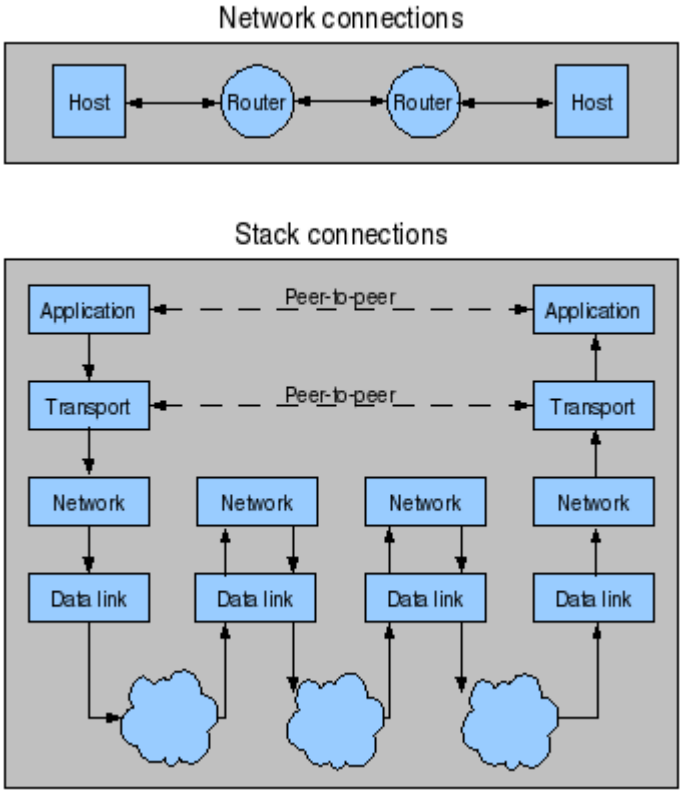

Fig. 3 - IP suite stack showing the physical network connection of two hosts via two routers and the corresponding layers used at each peer.

This abstraction also allows upper layers to provide services that the lower layers cannot, or choose not, to provide. For example, IP is designed to not be reliable and is a best effort delivery protocol. This means that all transport layer must address whether or not to provide reliability and to what degree. UDP provides data integrity (via a checksum) but does not guarantee delivery; TCP provides both data integrity and delivery guarantee (by retransmitted until the receiver receives the packet).

Layer	Protocols
Application	DNS, TLS/SSL, TFTP, FTP, HTTP, IMAP, IRC, NNTP, POP3, SIP, SMTP, SNMP, SSH, TELNET, BitTorrent, RTP, rlogin, ENRP, ...
Transport	TCP, UDP, DCCP, SCTP, IL, RUDP, ...
Network	IP (IPv4, IPv6), ICMP, IGMP, ARP, RARP, ...
Link	Ethernet, Wi-Fi, Token ring, PPP, SLIP, FDDI, ATM, Frame Relay, SMDS, ...

Fig. 4 - Internet protocol suite.

4.2 Application layer

The application layer is the layer that most common network-aware programs work in with the aim of communicating across a network. Communication that occurs in this layer are application specific and data is passed from the program, in the format used internally by this application, and is encapsulated into a transport layer protocol.

The actual data sent over the network is passed into the application layer where it is encapsulated into the application layer protocol. From there, the data is passed down into the lower layer protocol in the transport layer.

The two most common lower layer protocols are TCP and UDP. Both of which require a port in order to use their service and most well-used applications have specific ports assigned to them (HTTP has port 80; FTP has port 21; etc.) for servers while clients use ephemeral ports.

4.3 Transport layer

The transport layer is the layer between the application and the network protocol (i.e., IP) and primarily provides the service of connecting applications together through the use of ports. Since IP provides only a best effort delivery, the transport layer is the first layer to address reliability.

For example, TCP is a connection-oriented protocol that addresses numerous reliability issues to provide a reliable byte stream:

- data arrives in-order
- data is has minimal error-correctness
- duplicate data is discarded
- lost/discarded packets are resent
- includes traffic congestion control

UDP is a connectionless datagram protocol. Like IP, it is a best effort or "unreliable" protocol. The only reliability issue that it addresses is error-correctness of the data (albiet through a weak checksum algorithm). UDP is typically used for applications such as streaming media (audio and video, etc) where on-time arrival is more important than reliability, or for simple query/response applications like DNS lookups, where the overhead of setting up a reliable connection is disproportionately large.

4.4 Network layer

As originally defined, the Network layer solves the problem of getting packets across a single network. Examples of such protocols are X.25, and the ARPANET's Host/IMP Protocol.

With the advent of the concept of Internetworking, additional functionality was added to this layer, namely getting data from the source network to the destination network. This generally involves routing the packet across a network of networks, known as an internet.

In the internet protocol suite, IP performs the basic task of getting packets of data from source to destination. IP can carry data for a number of different upper layer protocols; these protocols are each identified by a unique protocol number: ICMP and IGMP are protocols 1 and 2, respectively.

4.5 Link layer

The link layer is not really part of the internet protocol suite, but is the method used to pass packets from the network layer on two different hosts. This process can be controlled both in the software device driver for the network card, as well as on firmware or specialist chipsets. These will perform data link functions such as adding a packet header to prepare it for transmission, then actually transmit the frame over a physical medium.

5 INTRANET

An intranet is a private network that uses Internet Protocols, network connectivity, and possibly the public telecommunication system to securely share part of an organization's information or operations with its employees. Sometimes the term refers only to the most visible service, the internal website.

The same concepts and technologies of the Internet such as clients and servers running on the Internet protocol suite are used to build an intranet. HTTP and other Internet protocols are commonly used as well, especially FTP and email. There is often an attempt to use Internet technologies to provide new interfaces with corporate 'legacy' data and information systems

There does not necessarily have to be any access from the organization's internal network to the Internet itself. Where there is, there will usually be a firewall with a gateway through which all access takes place, along with user authentication, encryption of messages, and the use of virtual private networks (VPNs) that tunnel through the public network. Through such devices, company information and computing resources can be shared by employees working from external locations.

Increasingly, intranets are being used to deliver tools and applications (e.g. collaboration to facilitate working in groups and for teleconferences or sophisticated corporate directories, sales and CRM tools, project management, etc.) to advance productivity.

When part of an intranet is made accessible to customers, partners, suppliers, or others outside the company, that part becomes part of an *extranet*.

5.1 Advantages of an intranet

1. *Workforce productivity*: Intranets can help employees to quickly find and view information and applications relevant to their roles and responsibilities. Via a simple-to-use web browser interface, users can access data held in any database the organization wants to make available, anytime and - subject to security provisions - from anywhere, increasing employees' ability to perform their jobs faster, more accurately, and with confidence that they have the right information.
2. *Time*: With intranets, organisations can make more information available to employees on a "pull" basis (ie: employees can link to relevant information at a time which suits them) rather than being deluged indiscriminately by emails.
3. *Communication*: Intranets can serve as powerful tools for communication within an organisation, vertically and horizontally.

6 EXTRANET

An extranet is a private network that uses Internet protocols, network connectivity, and possibly the public telecommunication system to securely share part of a business's information or operations with suppliers, vendors, partners, customers or other businesses. An extranet can be viewed as part of a company's Intranet that is extended to users outside the company (eg: normally over the Internet). It has also been described as a "state of mind" in which the Internet is perceived as a way to do business with other companies as well as to sell products to customers.

An argument has been made that "extranet" is just a buzzword[1] for describing what institutions have been doing for decades, that is, interconnecting to each other to create private networks for sharing information.

An extranet requires *security* and *privacy*. These can include *firewalls*, server management, the issuance and use of *digital certificates* or similar means of user *authentication*, *encryption* of messages, and the use of *virtual private networks (VPNs)* that tunnel through the public network.

During the late 1990s and early 2000s, several industries started to use the term "extranet" to describe central repositories of shared data made accessible via the web only to authorised members of particular work groups.

For example, in the construction industry, project teams could login to and access a 'project extranet' to share drawings and documents, make comments, issue requests for information, etc. In 2003 in the United Kingdom, several of the leading vendors formed the Network of Construction Collaboration Technology Providers, or NCCTP, to promote the technologies and to establish data exchange standards between the different systems. The same type of construction-focused technologies have also been developed in the United States, Australia, Scandinavia, Germany and Belgium, among others. Some applications are offered on a Software as a Service (SaaS) basis by vendors functioning as Application service providers (ASPs).

Specially secured extranets are used to provide virtual data room services to companies in several sectors (including law and accountancy).

There are a variety of commercial extranet applications, some of which are for pure file management, and others which include broader collaboration and project management tools.

7 GEOGRAPHICAL NETWORKS

7.1 Local Area Network (LAN)

A **local area network** (LAN) is a computer network covering a small local area, like a home, office, or small group of buildings such as a home, office, or college. Current LANs are most likely to be based on switched *Ethernet* or *Wi-Fi* technology running at from 10 to 10000

[1] A *buzzword* (also known as a fashion word or vogue word) is an idiom, often a neologism, commonly used in managerial, technical, administrative, and sometimes political environments. Buzzwords appear ubiquitously but their actual meanings often remain unclear. A buzzword may or may not appear in a dictionary, and if it does, its meaning as a buzzword may not match the conventional definition. (See http://en.wikipedia.org/wiki/Buzzword and http://en.wikipedia.org/wiki/List_of_buzzwords for more details).

Mbit/s. The defining characteristics of LANs in contrast to WANs are: a) much higher data rates, b) smaller geographic range - at most a few kilometers - and c) they do not involve leased telecommunication lines.

7.2 Wide Area Network (WAN)

A wide area network or WAN is a computer network covering a wide geographical area, involving a vast array of computers. This is different from personal area networks (PANs), metropolitan area networks (MANs) or local area networks (LANs) that are usually limited to a room, building or campus. The most well-known example of a WAN is the Internet.

WANs are used to connect local area networks (LANs) together, so that users and computers in one location can communicate with users and computers in other locations. Many WANs are built for one particular organization and are private. Others, built by Internet service providers, provide connections from an organization's LAN to the Internet.

7.3 Personal Area Network (PAN)

A personal area network (PAN) is a computer network used for communication among computer devices (including telephones and personal digital assistants) close to one person. The devices may or may not belong to the person in question. The reach of a PAN is typically *a few meters*.

PANs can be used for communication among the personal devices themselves (intrapersonal communication), or for connecting to a higher level network and the Internet (an uplink).

Personal area networks may be wired with computer buses such as USB and FireWire. A wireless personal area network (WPAN) can also be made possible with network technologies such as IrDA and Bluetooth.

7.4 Metropolitan Area Network (MAN)

Metropolitan Area Networks or MANs are large computer networks usually spanning a campus or a city. They typically use wireless infrastructure or optical fiber connections to link their sites.

For instance a university or college may have a MAN that joins together many of their local area networks (LANs) situated around site of a fraction of a square kilometer. Then from their MAN they could have several wide area network (WAN) links to other universities or the Internet.

8 Wireless Networks

8.1 Wireless LAN or WLAN

A **wireless LAN** or **WLAN** is a *wireless local area network* that uses *radio waves* as its carrier: the last link with the users is wireless, to give a network connection to all users in the surrounding area.

Areas may range from a single room to an entire campus. The backbone network usually uses cables, with one or more wireless access points connecting the wireless users to the wired network.

WLAN is expected to continue to be an important form of connection in many business areas. The market is expected to grow as the benefits of WLAN are recognized. So far WLANs have been installed in universities, airports, and other major public places. Decreasing costs of WLAN equipment has also brought it to many homes. Large future markets are estimated to be in health care, corporate offices and the downtown area of major cities. New York City has even begun a pilot program to cover all five boroughs of the city with wireless internet.

Originally WLAN hardware was so expensive that it was only used as an alternative to cabled LAN in places where cabling was difficult or impossible. Such places could be old protected buildings or classrooms, although the restricted range of the 802.11b (typically 30ft.) limits its use to smaller buildings. WLAN components are now cheap enough to be used in the home, with many being set-up so that one PC (a parent's PC, for example) can be used to share an internet connection with the whole family (whilst retaining access control at the parents' PC).

Early development included industry-specific solutions and proprietary protocols, but at the end of the 1990s these were replaced by standards, primarily the various versions of *IEEE 802.11 (Wi-Fi)* and *HomeRF* (2 Mbit/s, intended for home use, unknown in the UK). Recently the faster 54 Mbit/s 802.11a (5 GHz) and 802.11g (2.4 GHz) standards were released.

8.2 Wi-Fi

Wi-Fi (also WiFi, Wi-fi, Wifi, or wifi) is a set of product compatibility standards for wireless local area networks (WLAN) based on the IEEE 802.11 specifications. New standards beyond the 802.11 specifications, such as 802.16(WiMAX), are currently in the works and offer many enhancements, anywhere from longer range to greater transfer speeds.

Wi-Fi was intended to be used for mobile devices and LANs, but is now often used for Internet access. It enables a person with a wireless-enabled computer or personal digital assistant (PDA) to connect to the Internet when in proximity of an access point. The geographical region covered by one or several access points is called a *hotspot*.

Contrary to popular belief, Wi-Fi did not originally stand for Wireless-Fidelity. The term "Wi-Fi" was developed by the Wi-Fi Alliance along with the Interbrand Corporation (here) to describe WLAN products that are based on the IEEE 802.11 standards. Phil Belanger of the Wi-Fi Alliance quoted, "*Wi-Fi and the yin yang style logo were invented by Interbrand. We (the founding members of the Wireless Ethernet Compatibility Alliance, now called the Wi-Fi Alliance) hired Interbrand to come up with the name and logo that we could use for our interoperability seal and marketing efforts. We needed something that was a little catchier than "IEEE 802.11b Direct Sequence*". Later, the term "Wireless Fidelity" was coined with the marketing of a new tag line, "The Standard for Wireless Fidelity." But that was soon dropped due to confusion among customers and consumers.

8.3 IEEE 802.11

IEEE 802.11, the Wi-Fi standard, denotes a set of Wireless LAN/WLAN standards developed by working group 11 of the IEEE LAN/MAN Standards Committee (IEEE 802).

The term 802.11x is also used to denote this set of standards, and is not to be mistaken for any

one of its elements. There is no single 802.11x standard. The term IEEE 802.11 is also used to refer to the original 802.11, which is now sometimes called "802.11legacy."

8.4 Wireless access point

In computer networking, a **wireless access point (WAP or AP)** is a device that connects wireless communication devices together to form a wireless network. The WAP usually connects to a wired network, and can relay data between wireless devices and wired devices.

Fig. 5 - Wireless access point Planet WAP-4000.

8.5 Hotspots

Hotspots are locations where you can have access from mobile computers (such as a laptop or a PDA) without connection cables to networked services such as the internet. Hotspots are often found near restaurants, train stations, airports, cafes, libraries and other public places.

Most Wi-Fi hotspot equipment is IEEE 802.11b or IEEE 802.11g compliant and offers some level of security like WEP and/or WPA.

Many laptops come with IEEE 802.11b or IEEE 802.11g adapters built in. However, it is also possible to buy PCMCIA or USB based external adapters.

9 Virtual Private Network (VPN)

A **Virtual Private Network**, or VPN, is a *private communications network* usually used within a company, or by several different companies or organizations, to communicate over a public network.

VPN message traffic is carried on public networking infrastructure (e.g. the Internet) using standard (often insecure) protocols, or over a service provider's network providing VPN service guarded by well defined *Service Level Agreement (SLA)* between the VPN customer and the VPN service provider.

VPN involves two parts: the protected or "inside" network that provides physical security and administrative security sufficing to protect transmission (somehow, it is not always the case), and a less trust-worthy or "outside" network or segment (Internet is the largest "jungle"). Generally, a firewall sits between a remote user's workstation or client and the host network or server. As the user's client establishes the communication with the firewall, the client may pass authentication data to an authentication service inside the perimeter. A known trusted person, sometimes only when using trusted devices, can be provided with appropriate security

privileges to access resources not available to general users.

Many VPN client programs can be configured to require that all IP traffic must pass through the tunnel while the VPN is active, for better security. From the user's perspective, this means that while the VPN client is active, all access outside their employer's secure network must pass through the same firewall as would be the case while physically connected to the office ethernet. This reduces the risk that an attacker might gain access to the secured network by attacking the employee's laptop: to other computers on the employee's home network, or on the public internet, it is as though the machine running the VPN client simply does not exist. Such security is important because other computers local to the network on which the client computer is operating may be untrusted or partially trusted. Even with a home network that is protected from the outside internet by a firewall, people who share a home may be simultaneously working for different employers over their respective VPN connections from the shared home network. Each employer would therefore want to ensure their proprietary data is kept secure, even if another computer in the local network gets infected with malware. And if a travelling employee uses a VPN client from a Wi-Fi access point in a public place, such security is even more important.

Secure VPNs (SVPNs) use cryptographic tunneling protocols to provide the necessary confidentiality (preventing snooping), sender authentication (preventing identity spoofing), and message integrity (preventing message alteration) to achieve the privacy intended. When properly chosen, implemented, and used, such techniques can provide secure communications over unsecured networks. Because such choice, implementation, and use are not trivial, there are many insecure VPN schemes on the market.

Trusted VPNs do not use cryptographic tunneling, and instead rely on the security of a single provider's network to protect the traffic.

9.1 Types of VPN

The VPN market has tremendously expanded these last years. As it evolves the lines between various types of VPN classifications and architectures blur out. Hardware manufacturers now provide software clients that offer features historically available only through software or firewall-based solutions, while stand-alone applications on the other hand may support encrypting routers to improve performance. From a broad standpoint, one can identify three basic types of VPN.

- **Intranet VPN**: this type of VPN is client transparent. Intranet VPN is usually implemented for networks within a common network infrastructure but across various physical locations. For instance several buildings may be connected to a data center, or a common mainframe application that they can access securely through private lines. Intranet VPNs need to be especially secure with strong encryption and meet strict performance and bandwidth requirements. They must remain easily upgradeable since many VPN clients or users may be added (additional locations or applications).
- **Remote Access VPN**: this type of VPN is client initiated and intended for salesmen equipped with laptops and telecommuters that will connect intermittently from vary diverse locations (homes, hotels, conference halls...). The key factor of remote access VPN is flexibility as performance and bandwidth are usually minimal and less of an issue. More than encryption, authentication will be the main security concern for remote access VPN.

- **Extranet VPN**: in this case VPN uses the Internet as main backbone. Extranet VPN usually addresses a wider scale of users and locations, enabling customers, suppliers and branch offices to access corporate resources across various network architectures.

10 REFERENCES

[1] Wikipedia: http://en.wikipedia.org

[2] Wilkinson, Paul (2005). *Collaboration Technologies: The Extranet Evolution*, Taylor & Francis.