

7. Communications protocols

The TCP/IP model or Internet layering scheme and its relation to common protocols often layered on top of it.

A communications protocol is a set of rules for exchanging information over network links. In a protocol stack (also see the OSI model), each protocol leverages the services of the protocol below it. An important example of a protocol stack is HTTP running over TCP over IP over IEEE 802.11. (TCP and IP are members of the Internet Protocol Suite. IEEE 802.11 is a member of the Ethernet protocol suite.) This stack is used between the wireless router and the home user's personal computer when the user is surfing the web.

Whilst the use of protocol layering is today ubiquitous across the field of computer networking, it has been historically criticized by many researchers for two principal reasons. Firstly, abstracting the protocol stack in this way may cause a higher layer to duplicate functionality of a lower layer, a prime example being error recovery on both a per-link basis and an end-to-end basis. Secondly, it is common that a protocol implementation at one layer may require data, state or addressing information that is only present at another layer, thus defeating the point of separating the layers in the first place. For example, TCP uses the ECN field in the IPv4 header as an indication of congestion; IP is a network layer protocol whereas TCP is a transport layer protocol.

Communication protocols have various characteristics. They may be connection-oriented or connectionless, they may use circuit mode or packet switching, and they may use hierarchical addressing or flat addressing.

There are many communication protocols, a few of which are described below.

Ethernet

Ethernet is a family of protocols used in LANs, described by a set of standards together called IEEE 802 published by the Institute of Electrical and Electronics Engineers. It has a flat addressing scheme. It operates mostly at levels 1 and 2 of the OSI model. For home users today, the most well-known member of this protocol family is IEEE 802.11, otherwise known as Wireless LAN (WLAN). The complete IEEE 802 protocol suite provides a diverse set of networking capabilities. For example, MAC bridging (IEEE 802.1D) deals with the routing of Ethernet packets using a Spanning Tree Protocol, IEEE 802.1Q describes VLANs, and IEEE 802.1X defines a port-based Network Access Control protocol, which forms the basis for the

authentication mechanisms used in VLANs (but it is also found in WLANs) – it is what the home user sees when the user has to enter a "wireless access key".

7.1 Internet Protocol Suite

The Internet Protocol Suite, also called TCP/IP, is the foundation of all modern networking. It offers connection-less as well as connection-oriented services over an inherently unreliable network traversed by data-gram transmission at the Internet protocol (IP) level. At its core, the protocol suite defines the addressing, identification, and routing specifications for Internet Protocol Version 4 (IPv4) and for IPv6, the next generation of the protocol with a much enlarged addressing capability.

7.2 SONET/SDH

Synchronous optical networking (SONET) and Synchronous Digital Hierarchy (SDH) are standardized multiplexing protocols that transfer multiple digital bit streams over optical fiber using lasers. They were originally designed to transport circuit mode communications from a variety of different sources, primarily to support real-time, uncompressed, circuit-switched voice encoded in PCM (Pulse-Code Modulation) format. However, due to its protocol neutrality and transport-oriented features, SONET/SDH also was the obvious choice for transporting Asynchronous Transfer Mode (ATM) frames.

7.3 Asynchronous Transfer Mode

Asynchronous Transfer Mode (ATM) is a switching technique for telecommunication networks. It uses asynchronous time-division multiplexing and encodes data into small, fixed-sized cells. This differs from other protocols such as the Internet Protocol Suite or Ethernet that use variable sized packets or frames. ATM has similarity with both circuit and packet switched networking. This makes it a good choice for a network that must handle both traditional high-throughput data traffic, and real-time, low-latency content such as voice and video. ATM uses a connection-oriented model in which a virtual circuit must be established between two endpoints before the actual data exchange begins.

While the role of ATM is diminishing in favor of next-generation networks, it still plays a role in the last mile, which is the connection between an Internet service provider and the home user. For an interesting write-up of the technologies involved, including the deep stacking of communications protocols used, see.

7.4 Geographic scale

A network can be characterized by its physical capacity or its organizational purpose. Use of the network, including user authorization and access rights, differ accordingly.

7.4.1 Personal area network

A personal area network (PAN) is a computer network used for communication among computer and different information technological devices close to one person. Some examples of devices that are used in a PAN are personal computers, printers, fax machines, telephones, PDAs, scanners, and even video game consoles. A PAN may include wired and wireless devices. The reach of a PAN typically extends to 10 meters. A wired PAN is usually constructed with USB and FireWire connections while technologies such as Bluetooth and infrared communication typically form a wireless PAN.

7.4.2 Local area network

A local area network (LAN) is a network that connects computers and devices in a limited geographical area such as a home, school, office building, or closely positioned group of buildings. Each computer or device on the network is a node. Wired LANs are most likely based on Ethernet technology. Newer standards such as ITU-T G.hn also provide a way to create a wired LAN using existing wiring, such as coaxial cables, telephone lines, and power lines.

All interconnected devices use the network layer (layer 3) to handle multiple subnets (represented by different colors). Those inside the library have 10/100 Mbit/s Ethernet connections to the user device and a Gigabit Ethernet connection to the central router. They could be called Layer 3 switches, because they only have Ethernet interfaces and support the Internet Protocol. It might be more correct to call them access routers, where the router at the top is a distribution router that connects to the Internet and to the academic networks' customer access routers.

The defining characteristics of a LAN, in contrast to a wide area network (WAN), include higher data transfer rates, limited geographic range, and lack of reliance on leased lines to provide connectivity. Current Ethernet or other IEEE 802.3 LAN technologies operate at data transfer rates up to 10 Gbit/s. The IEEE investigates the standardization of 40 and 100 Gbit/s rates. A LAN can be connected to a WAN using a router.

7.4.3 Home area network

A home area network (HAN) is a residential LAN used for communication between digital devices typically deployed in the home, usually a small number of personal computers and accessories, such as printers and mobile computing devices. An important function is the sharing of Internet access, often a broadband service through a cable TV or digital subscriber line (DSL) provider.

7.4.4 Storage area network

A storage area network (SAN) is a dedicated network that provides access to consolidated, block level data storage. SANs are primarily used to make storage devices, such as disk arrays, tape libraries, and optical jukeboxes, accessible to servers so that the devices appear like locally attached devices to the operating system. A SAN typically has its own network of storage devices that are generally not accessible through the local area network by other devices. The cost and complexity of SANs dropped in the early 2000s to levels allowing wider adoption across both enterprise and small to medium sized business environments.

7.4.5 Campus area network

A campus area network (CAN) is made up of an interconnection of LANs within a limited geographical area. The networking equipment (switches, routers) and transmission media (optical fiber, copper plant, Cat5 cabling, etc.) are almost entirely owned by the campus tenant / owner (an enterprise, university, government, etc.).

For example, a university campus network is likely to link a variety of campus buildings to connect academic colleges or departments, the library, and student residence halls.

7.4.6 Backbone network

A backbone network is part of a computer network infrastructure that provides a path for the exchange of information between different LANs or sub-networks. A backbone can tie together diverse networks within the same building, across different buildings, or over a wide area.

For example, a large company might implement a backbone network to connect departments that are located around the world. The equipment that ties together the departmental networks constitutes the network backbone. When designing a network backbone, network performance and network congestion are critical factors to take

into account. Normally, the backbone network's capacity is greater than that of the individual networks connected to it.

Another example of a backbone network is the Internet backbone, which is the set of wide area networks (WANs) and core routers that tie together all networks connected to the Internet.

7.4.7 Metropolitan area network

A Metropolitan area network (MAN) is a large computer network that usually spans a city or a large campus.

7.4.8 Wide area network

A wide area network (WAN) is a computer network that covers a large geographic area such as a city, country, or spans even intercontinental distances. A WAN uses a communications channel that combines many types of media such as telephone lines, cables, and air waves. A WAN often makes use of transmission facilities provided by common carriers, such as telephone companies. WAN technologies generally function at the lower three layers of the OSI reference model: the physical layer, the data link layer, and the network layer.

7.4.9 Enterprise private network

An enterprise private network is a network that a single organization builds to interconnect its office locations (e.g., production sites, head offices, remote offices, shops) so they can share computer resources.

7.4.10 Virtual private network

A virtual private network (VPN) is an overlay network in which some of the links between nodes are carried by open connections or virtual circuits in some larger network (e.g., the Internet) instead of by physical wires. The data link layer protocols of the virtual network are said to be tunneled through the larger network when this is the case. One common application is secure communications through the public Internet, but a VPN need not have explicit security features, such as authentication or content encryption. VPNs, for example, can be used to separate the traffic of different user communities over an underlying network with strong security features.

VPN may have best-effort performance, or may have a defined service level agreement (SLA) between the VPN customer and the VPN service provider. Generally, a VPN has a topology more complex than point-to-point.

7.4.11 Global area network

A global area network (GAN) is a network used for supporting mobile across an arbitrary number of wireless LANs, satellite coverage areas, etc. The key challenge in mobile communications is handing off user communications from one local coverage area to the next. In IEEE Project 802, this involves a succession of terrestrial wireless LANs.

7.4.12 Organizational scope

Networks are typically managed by the organizations that own them. Private enterprise networks may use a combination of intranets and extranets. They may also provide network access to the Internet, which has no single owner and permits virtually unlimited global connectivity.

7.5 Intranets

An intranet is a set of networks that are under the control of a single administrative entity. The intranet uses the IP protocol and IP-based tools such as web browsers and file transfer applications. The administrative entity limits use of the intranet to its authorized users. Most commonly, an intranet is the internal LAN of an organization. A large intranet typically has at least one web server to provide users with organizational information. An intranet is also anything behind the router on a local area network.

7.6 Extranet

An extranet is a network that is also under the administrative control of a single organization, but supports a limited connection to a specific external network. For example, an organization may provide access to some aspects of its intranet to share data with its business partners or customers. These other entities are not necessarily trusted from a security standpoint. Network connection to an extranet is often, but not always, implemented via WAN technology.

7.7 Internetwork

An internetwork is the connection of multiple computer networks via a common routing technology using routers.

7.8 Internet

Partial map of the Internet based on the January 15, 2005 data found on opte.org. Each line is drawn between two nodes, representing two IP addresses. The length of

the lines are indicative of the delay between those two nodes. This graph represents less than 30% of the Class C networks reachable.

The Internet is the largest example of an internetwork. It is a global system of interconnected governmental, academic, corporate, public, and private computer networks. It is based on the networking technologies of the Internet Protocol Suite. It is the successor of the Advanced Research Projects Agency Network (ARPANET) developed by DARPA of the United States Department of Defense. The Internet is also the communications backbone underlying the World Wide Web (WWW).

Participants in the Internet use a diverse array of methods of several hundred documented, and often standardized, protocols compatible with the Internet Protocol Suite and an addressing system (IP addresses) administered by the Internet Assigned Numbers Authority and address registries. Service providers and large enterprises exchange information about the reachability of their address spaces through the Border Gateway Protocol (BGP), forming a redundant worldwide mesh of transmission paths.

7.10 Darknet

A Darknet is an overlay network, typically running on the internet, that is only accessible through specialized software. A darknet is an anonymizing network where connections are made only between trusted peers — sometimes called "friends" (F2F) — using non-standard protocols and ports.

Darknets are distinct from other distributed peer-to-peer networks as sharing is anonymous (that is, IP addresses are not publicly shared), and therefore users can communicate with little fear of governmental or corporate interference.

7.11 Routing

Routing calculates good paths through a network for information to take. For example from node 1 to node 6 the best routes are likely to be 1-8-7-6 or 1-8-10-6, as this has the thickest routes.

Routing is the process of selecting network paths to carry network traffic. Routing is performed for many kinds of networks, including circuit switching networks and packet switched networks.

In packet switched networks, routing directs packet forwarding (the transit of logically addressed network packets from their source toward their ultimate destination) through intermediate nodes. Intermediate nodes are typically network hardware devices such as routers, bridges, gateways, firewalls, or switches. General-purpose computers can also forward packets and perform routing, though they are not specialized hardware and may suffer from limited performance. The routing process usually directs forwarding on the basis of routing tables, which maintain a record of the routes to various network destinations. Thus, constructing routing tables, which are held in the router's memory, is very important for efficient routing. Most routing algorithms use only one network path at a time. Multipath routing techniques enable the use of multiple alternative paths.

There are usually multiple routes that can be taken, and to choose between them, different elements can be considered to decide which routes get installed into the routing table, such as (sorted by priority):

Prefix-Length: where longer subnet masks are preferred (independent if it is within a routing protocol or over different routing protocol)

Metric: where a lower metric/cost is preferred (only valid within one and the same routing protocol)

Administrative distance: where a lower distance is preferred (only valid between different routing protocols)

Routing, in a more narrow sense of the term, is often contrasted with bridging in its assumption that network addresses are structured and that similar addresses imply proximity within the network. Structured addresses allow a single routing table entry to represent the route to a group of devices. In large networks, structured addressing (routing, in the narrow sense) outperforms unstructured addressing (bridging). Routing has become the dominant form of addressing on the Internet. Bridging is still widely used within localized environments.

7.12 Network service

Network services are applications hosted by servers on a computer network, to provide some functionality for members or users of the network, or to help the network itself to operate.

The World Wide Web, E-mail, printing and network file sharing are examples of well-known network services. Network services such as DNS (Domain Name

System) give names for IP and MAC addresses (people remember names like “nm.lan” better than numbers like “210.121.67.18”), and DHCP to ensure that the equipment on the network has a valid IP address.

Services are usually based on a service protocol that defines the format and sequencing of messages between clients and servers of that network service.