A Busy Person's Guide to How the Internet Works (and is paid for)

Short Summary | August 2022

Overview:

The Internet is a collection of people, nodes, networks and content. Often it is described as a "cloud" made up of connections between devices, but this is misleading because there are rarely direct connections over the Internet. In fact, the Internet is completely decentralised with many nodes and combinations of direct and indirect connections between them.

How Do You Get Access to the Internet?

Consumers and enterprises both gain access to the Internet through Internet Service Providers (ISPs). ISPs sell access to the Internet as a service. For a consumer, this often means a fixed Internet access at home or Internet access for a mobile device. In the first case the ISP often provides a router or home hub that allows devices in the home to connect. Messages from nodes inside the home are directed to the router and then out to the Internet through the ISPs network. Enterprises and businesses usually have connections with more capacity. Still, the basic picture is often the same: the enterprise has a router on the "edge" of its internal network and nodes inside the network send messages through that router and out to the ISP's network connections.

Many years ago, most connections used the copper wires that also provided traditional voice communications. Today, the Internet uses a variety of access technologies that allow people and businesses to connect, including cable services, cellular and mobile Internet, fibre optic, fixed wireless and even access to the Internet via satellites.

ISPs get paid by end-users (either consumers or enterprises) for Internet access provision. End-users are generally charged on an unmetered basis through "flat rates" including the Internet connection itself and the transfer of an unlimited volume of data.

What Does Information Look Like on the Internet?

The basic unit of information on the Internet is called a "packet." Every time a node sends messages on the Internet, it constructs a series of packets that function as envelopes containing the information being exchanged. The packets usually have two parts: a "header" that contains information about the packet, and the "payload" which contains the content of the message.

Packets are simply made up of ones and zeroes, known as "bits." All messages traverse the Internet in this binary format. The nodes in the network are responsible for encoding information (email, pictures, video, etc.) from its original format into bits that can be assembled into packets. One of the most important parts of the header is

the addressing information it contains. Each packet's header has the address of the node the message is being sent to, as well as the address the message is being sent from. Routers in the Internet use these addresses to decide how to forward the packet on and closer to its intended destination.

Most things sent on the Internet are too big to fit in a single packet, so a disassembly and assembly process takes place for almost every service or application on the Internet. For example, an image can be made up of many packets, with the sending node breaking the image into multiple packets and sending all of them through the routers toward the receiving node. The destination would then take the packets and reassemble them into the original image.

It's crucial that both the sender and receiver use the same approach for disassembly and assembly of information. To ensure that the sender and receiver can understand each other, a set of standard, interoperable "protocols" have been developed so that all nodes will have a common set of rules for exchanging messages and information. These protocols are simply rules for nodes to use to talk to each other. Very famous protocols include the Hypertext Transport Protocol (HTTP), which allows the exchange of messages for the World Wide Web and the Simple Mail Transport Protocol (SMTP), which allows for the transport of messages for electronic mail.

Applications as complex as video conferencing and as simple as Internet chat all have this defining characteristic in common: the use of standardised protocols to exchange information contained in packets over the networks connected to the Internet.

How Does Information Move Around the Internet?

The Internet is not one unified network, it is a "network of networks." The networks that make up the building blocks of the Internet are called Autonomous Systems (AS), and there are about 100,000 independently administered ASes on the Internet today. When these networks interconnect, they constitute the public Internet as we know it.

To move information around the Internet there needs to be "routes" through the map of interconnected networks. The Border Gateway Protocol (BGP) allows computers to map a route for packets from source to destination using the shortest and cheapest path possible.

Each AS uses BGP to build its own map for traversing the Internet, but no BGP server has a complete global map of all the possible paths for packets on the Internet. When a node sends a packet onto the Internet, it usually sends the packet to the first router it knows about. That router reads the packet and looks at the header to decide if the packet is destined for a node that is connected to the same network as the router. If the packet is destined for a node on the ISP's own network, the ISP can deliver it on its own. If the address indicates that the packet must be sent onward, it uses BGP to decide which new router to send it to. Every router on the path does this until the packets arrive at their destination.

When BGP servers of different ASes talk to each other, they become neighbours. As neighbours, they exchange maps of the routes they know about and want to share. Shorter paths for the packets are preferred because that results in less time for the packet to get to its intended destination.

Once the ASes have a map of available routes with their neighbours, they can exchange data. In some cases, the ASes will agree to exchange traffic between the ASes for free – an arrangement called "peering." Often the size of the neighbours will be different. The word "transit" is used to describe connectivity to any destination on the Internet.

Any organisation responsible for part of a packet's journey from source to destination has a stake in peering and transit. This includes access providers of any size. ISPs who provide access – using any access technology – need

a way to route the packet toward its destination. In addition, the access provider also provides a path to its own customers. That way, packets headed to an ISP customer's endpoint have a path for delivery.

How is Peering and Transit Paid For?

There are agreements between the different actors so they can send their data to each other. Transit arrangements provide access to the entire Internet, while IP Peering arrangements facilitate the direct mutual exchange of traffic between connected players and each party's downstream customers. The two types of interconnections can be both complementary and substitutable arrangements depending on the network configuration chosen by an operator.

A key difference between transit and peering is that in the vast majority of cases, no money is exchanged between networks that agree to peer with one another. The benefit to an operator of an AS to peer is that the arrangement can reduce dependency on upstream transit providers – reducing that dependency has the effect of reducing cost for that operator.

Peering enables the two networks to exchange data and benefit equally. The peering arrangement also has the ASes advertise only their internal customer routes. Although this is a zero-cost arrangement, there are usually costs for both parties associated with co-location, the routes to get to the co-located facility, and the required infrastructure connections. Strategies for peering (and interconnection in general) vary from one network operator to another. Those may be explained in a reference document, which is public most of the time and known as the "peering policy", although there is rarely any requirement to stick to this policy in the network's dealings with others.

In practice, peering agreements are often not covered by a written contract and are established by informal agreements between the two parties. According to Packet Clearing House (PCH), in 2016 about 99.9 per cent of the peering agreements were done informally using a handshake.