History of the Internet Tutorial

Thanks for downloading the printable version of this tutorial. The original is here:
http://www.techopedia.com/6/27861/internet/history-of-the-internet

Table of Contents

Introduction

The Internet is ubiquitous in every sense of the word. We use it to guide us through unfamiliar cities, send photos to far-off friends, learn new tasks, for entertainment and to answer a multitude of other needs that come up every day. Although the Internet hasn't been around for that long, our dependence on it will soon run so deep that we will tend to think of its existence as a given – much like another ubiquitous invention: the light bulb.

However, whereas most people can give you a textbook version of the creation of the light bulb and name its inventor (Thomas Edison), the origins of the Internet are a jumbled mess of myth, simultaneous discovery and uncertain time lines. If something does come to mind, it is usually Al Gore's supposed claim to having invented the Internet. In this tutorial, we'll look at the origins of the Internet and the people who are really behind its creation.

What Is the Internet?

To tell the history of the Internet, you have to commit to a definition. The one that springs to mind most easily is a network for sharing information, but this definition is far too broad. Under this definition, a history of the Internet would start with the evolution of language, jump between continents to tell the story of tribal drumming, smoke signals, carrier pigeons, messengers, telegraphs, the transatlantic connection, phone lines and so on. (For background reading, check out What is the difference between the Internet and the World Wide Web?)

Similarly, adding the word "computer" to the definition helps eliminate a brief history of selective pigeon breeding, but it also limits the story a little too much. Computers are still a common interface that people use to connect to the Internet, but smartphones, tablets and future devices make for a laundry list definition where the Internet is a network for sharing information using personal computers, smartphones, e-readers and so on.

A better definition of the Internet, therefore, focuses on the method by which the information is being shared, regardless of the technology. And, at its heart, the tale of the Internet is really a tale of protocols, where protocol is loosely defined as "rules of communication". For example, the early telegraph system was a communication system where the protocol was the rules behind the dots and dashes telegraph operators used to understand the message.

The Federal Networking Council (FNC) recognized the importance of protocols when it released an official definition of the term on October 24, 1995. According to the FNC, the Internet is a global information system that meets the following three criteria:

- Is logically linked together by a globally unique address space based on Internet Protocol (IP) or its subsequent extensions/follow-ons
- Is able to support communications using the Transmission Control Protocol/Internet Protocol (TCP/IP) suite or its subsequent extensions/follow-ons, and/or other IP-compatible protocols
- Provides, uses or makes accessible, either publicly or privately, high-level services layered on the communications and related infrastructure described herein

Hence, to understand the history of the Internet, is to understand the development of its underlying protocols. Using this approach, now that we have a target, we can track the creation of this global information system from its theoretical roots.

Better Dead Than Red: The Theory and Motivations for the Internet

From the end of WWII through to the 1980s, fear of the Soviet Union was a powerful force that motivated nations to put more money into weapons and technology research. The United States, considered the U.S.S.R.'s primary opponent in the Cold War, was shocked when the U.S.S.R. launched a satellite called Sputnik on October 4, 1957. It was the first satellite ever to be put into orbit, and it heralded the start of a space race that would see both animals and men launched beyond gravity's pull.

ARPA Founded to Avoid Future Sputniks

In 1958, President Eisenhower created the [Advanced Research Projects Agency](#) (ARPA) in response to the Sputnik launch. ARPA was a place for the Department of Defense to focus on research and development projects so that the U.S. could avoid the unpleasant surprise of another Sputnik – that is, another nation demonstrating technological capabilities that the U.S. could not match.

One of the more terrifying aspects of Sputnik's launch was the realization that Soviet rocketry had evolved to the point where missiles could strike the United States. This put the entire U.S. infrastructure at risk, including its information infrastructure.

The Survival Network and the Need for Packet-Switching

Prior to Sputnik, the best way to protect vital information was to store it all in a centralized and heavily guarded location; the new possibility that that centralized location could be hit by a nuclear missile forced the Department of Defense to change its policy.

Instead of holding all the sensitive information in one location that could be targeted and destroyed, the U.S. began exploring ways to create a communications and storage network that could survive a strike.

A quick point has to be made here. It is a common misconception that the possibility of nuclear strike was the driving force behind the Internet. It wasn't. As mentioned, the launch of Sputnik spurred the creation of ARPA, which is the department that created ARPANET. However, by the time they created ARPANET, their primary motivation was to share computer resources. That said, the threat of nuclear war was the primary reason that the RAND corporation - and specifically Paul Baran - came up with the theoretical concept of distributed networks.

In order to have an effective network that can survive a nuclear strike, data needs to be broken into small pieces that can be sent rapidly and reassembled with no errors where there is not a single point of failure. In short, you need [packet switching](#). Like many conceptual inventions, its

creation is shared between three people at three different organizations.

Paul Baran and Nuke-Proof Network

In 1959, Paul Baran was given the intellectual challenge of creating a communications system that would operate even if intermediate points were destroyed in a nuclear strike. The problem with the existing system was that it was centralized and worked in a linear fashion: A message left point A, traveled to the central switching area and was sent to endpoint B. Destroying point A or point B only stopped communication in one area, but destroying a central switching area took down the entire communications system.

The answer Baran came up with was to share communication among multiple, interconnected points. Each point in the communication system was connected to several other points in a non-linear arrangement. Baran referred to these points as nodes and used them to create distributed networks. He ran simulations of nuclear attacks on these distributed networks and found they could sustain major damage and still maintain end-to-end communications. Moreover, the more connections there were per node, the more resilient the network became as redundancies were created.

Message Blocks and Packets

The key to getting information around this network was a routing doctrine known as "hot-potato" switching. Information would be divided into message blocks that were sent separately across the network. These blocks would be received by the nearest node and stored. That node would then look for the shortest path to send the message blocks to their destination node. If an area of the network was destroyed, the nodes would just send the message blocks around the missing areas.

The reason we don't talk about message block switching is that Donald Watts Davies came up with the same concept at around the same time as Baran. Davies worked for the National Physical Laboratory (NPL) in the United Kingdom. Faced with a similar problem – how to create a robust communications network – Davies came up with packet switching. The team at ARPA eventually built such a network by drawing from both Baran's and Davies' work. They chose to call the system packet switching.

Leonard Kleinrock's Research Project

Leonard Kleinrock was also working on packet-switching networks before he even knew what they were called. Kleinrock was doing theoretical research on waiting lines when he worked out the mathematical theory that would be the backbone of packet-switching networks. In fact,

Kleinrock got to see his theoretical work become reality when he participated in the ARPANET project – the first packet-switching network.

The Original Networks

ARPANET is probably the network to which the origins of the Internet are most strongly tied, but it is not the only one. There were several networks, each of which brought a unique function that we now consider integral to the modern Internet. In this section, we'll look at each network as well as its contribution to the whole. (For background reading, see A Timeline of the Development of the Internet and World Wide Web.)

ARPANET

Although the work done by Paul Baran was prompted by the threat of nuclear attack, ARPA used Baran's research to lay the foundation of a network designed to solve the issues it was having with limited computing resources.

ARPA had sponsored computers in different locations, each connected back to the ARPA offices. Unfortunately, to communicate with a particular location, the ARPA information processing director, Robert Taylor, would have to go to the terminal connected to that location. This meant switching terminals every time he wished to contact someone else.

The inconvenience of the setup prompted Taylor to search for a way to network the computers so that a single terminal could communicate with all the locations. To this end, he brought in Larry Roberts, who created a plan for a packet-switching network between the computers named ARPANET.

Armed with a plan, ARPA sent out a request for quotation (RFQ) for creating the necessary interface message processors (IMPs). This contract was won by Bolt Beranek and Newman (now BBN Technologies, a subsidiary of Raytheon). Other necessary components were brought together by teams at UCLA and the Network Analysis Corporation, among others.

By 1969, ARPANET was coming together using a system of requests for comment (RFCs) that helped the designers work out the issues as hosts were being connected to the network. The first host-to-host message was sent from UCLA to the Stanford Research Institute (SRI) and consisted of Kleinrock watching his student, Charley Kline, type while confirming over the phone that the letters were appearing on the screen in Stanford. The system crashed after the "g" in "login," but the first Internet message had been sent.

Other Important Networks

Although ARPANET plays a central role up to this point in Internet history and going forward, there were other networks that brought something to the Internet in the form of alternative methods, or even just confirmation of important concepts. We'll look at these briefly.

Mark I

Donald Watts Davies was not idle after independently coming up with packet switching. He created proposals for a national data network that would span the U.K., but he didn't have the support or the resources that ARPA had. Davies was, however, successful in creating a packet-switching network to serve the needs of the National Physical Laboratory. He created a small-scale network called Mark I in 1970, and then an improved version, Mark II, in 1973. The Mark II network was still operational in 1986, but its reach was limited by funding and resources.

Merit Network

Merit Network was another network created in near parallel with the ARPANET timeline. It was formed in 1966 with the intention of linking together Michigan's public university computers. It got off the ground in 1971 thanks to funding from the National Science Foundation (NSF). It evolved alongside ARPANET, and eventually became a cornerstone of the NSFNET, which served as a backbone for the modern Internet.

Cyclades

Cyclades was a French-led packet-switching network that was inspired by ARPANET, but was designed to test out alternatives. Started in 1972, Cyclades pioneered several technical innovations that were later incorporated back into the ARPANET design. These innovations included layered architecture, adaptive routing, and making the host, rather than the network, responsible for the data (end to end). These innovations reappeared as part of TCP/IP on ARPANET.

Back to ARPANET

ARPANET evolved rapidly after the initial nodes were added, and its RFC system became fundamental to the creation of the Internet protocols we still depend on today. Network Control Protocol (NCP) served as the original host-to-host protocol as more and more hosts were added in the early '70s.

It is, however, dishonest to pretend that the history of the Internet was centered just on ARPANET. In truth, the innovations that took the concept of a few interconnected computers

and transformed them into vast networks of networks came from many sources. There were the aforementioned networks as well as unmentioned organizations that played important roles. That said, the story is more easily told from an ARPANET-centric view - both for clarity and for two very important innovations that grew out of the project: TCP/IP and NSFNET.

TCP/IP: One Protocol to Rule Them All

As packet-switching networks began to pop up in different areas of the world, the desire to have them all communicate with one another increased. Working off of the ideas used to create Cyclades, Robert Kahn and Vinton Cerf began to work on a new protocol to replace NCP. Transmission Control Protocol (TCP) was being tested on ARPANET as an alternative in 1974.

TCP improved the reliability of data transfer across the network, but it was only half of the puzzle. Cerf and Kahn added Internet Protocol (IP) to form the complementary TCP/IP protocol that is still in use today. TCP/IP allowed for end-to-end connectivity and specified how data passing through the network would be:

- Formatted
- Addressed
- Transmitted
- Routed
- Received

Again, the RFC system was integral to refining these protocols, and TCP/IP borrowed many features from the protocols designed by other organizations.

With TCP/IP we began to see a new practice of collecting a group of other procotols – we will inaccurately call them sub-protocols for the purpose of this explanation – and group them together into a parent protocol, which is considered a standard unto itself.

Over this same period, ARPA was renamed the Defense Advanced Research Projects Agency(DARPA), and steps were taken to segment a military network from research-based ARPRANET. This military network was officially launched in 1983 under the name MILNET.

TCP/IP Connects the Dots

TCP/IP served as the Swiss army knife of protocols, allowing diverse machines and network hardware to connect to a common network. It was officially implemented by ARPANET in 1983. By this time, ARPANET was just one of many established networks, some of which were already using TCP/IP, while others were using X.25.

Throughout the U.S. and Europe, these packet-switching networks were adding more and more nodes. Major networks were used as backbone networks, forming cores that served as high-traffic routes for data to pass through on the way to its destination. In addition to Mark II, Cyclades, Telenet and so on, some of these networks included NSFNET, a network funded by the National Science Foundation (NSF) to link educational institutions and CSNET, a network tied into ARPANET over a dial-up connection using X.25.

After some competition between protocols, TCP/IP was chosen as a standard protocol, making it possible to start connecting the dots across the world. In the U.S., NSFNET played a vital role by providing many of the Internet backbones and helping establish federal Internet exchanges. Theses spelled the end of ARPANET as a central player in the story of the Internet.

Beginnings of the Commercial Internet

From the late 1970s on, the importance of the connectivity provided by the NSF and its educational alliances cannot be overstated. There was no commercial viability or the Internet at this time, but it continued to grow at an incredible rate. The reason was the community of universities that took it upon themselves to provide the hubs necessary for a large volume of interconnectivity. Originally, telecommunications budgets were justified by the research aspect of this new Internet, but the volume of traffic grew to the point where there was a significant commitment required to continue growth.

Without the support of universities and their dedication to the ideal of a "free" Internet, there simply would not have ever been a commercially viable Internet. This idea that the Internet was to be free – and everything on it should be free as well – was both the biggest single reason why the Internet became a commercial success and the biggest single problem to its commercial viability.

Throughout the 1980s, however, the technical and hardware barriers to a commercially available Internet were systematically being eliminated. Telecommunications companies were investing heavily in networking infrastructure, while companies like Xerox, Intel and DEC were developingEthernet, laying the groundwork for both wide and local area networks.

The Invention of Email

As more hosts were added, more people began working on protocols and applications for the Internet. The first of these applications to meet our modern understanding of a killer app was email. Prior to email, communications between computers were like a chat. This had limitations in that you could only talk to someone in real time.

In 1972, Ray Tomlinson, an ARPANET contractor, tackled this problem. He came up with the

addressing system – user@host computer – as well as the basics of an email message. Messages were sent using utilities called SNDMSG and read using READMAIL. The RFC system helped refine email, and soon email represented a majority of all the traffic going back and forth over the network.

ARPANET can also lay claim to the first instance of spam, when Gary Thuerk sent a mass email about a new machine at Digital Equipment Corp. (DEC) on May 1, 1978.

In a grand perspective, email was the first compelling reason for making the Internet commercially available. The ability to send electronic messages to any connected computer gained value as more and more computers were networked together. Even now, email is one of the primary ways in which people use their Internet connections.

Telenet and Usenet

Even though most of the network power in the world was being used for research or email, some people still found a way to have fun. Telenet and Usenet provided examples of how the Internet might one day become a tool for entertainment.

Usenet was launched in 1980 by two graduate students at Duke University. It was the first example of a public Internet discussion group that served as a precursor to bulletin boards and Web forums. Telenet was commercial – and one of the activities people were using the connection for was to play a text based role-playing game in the form of a multi-user dungeon (MUD), foreshadowing the massively multiplayer online role-playing games (MMORPG) that would become popular in the future. In both cases, however, the people using networks for entertainment had to have a high degree of technical knowledge just to get started.

Stumbling Around the Web - A Few More Languages and Protocols

By the 1980s, the Internet was a worldwide network, but it was still missing some essential parts that would enable non-technical people to use it. Before we get to the Web browser, there are a few more concepts that need to be understood.

The Web

The World Wide Web (WWW) and the Internet are technically different things, even though we use the terms interchangeably. The Internet was the quiet, hardworking stage hand that did all the hard work behind the scenes, while the WWW became the sexy, young thing that attracted all the audience's attention.The Web could not exist without the Internet, so we often lump the two together historically. However, the Web is actually much younger. It traces its roots to 1980

and Tim Berners-Lee, who was working at the European Organization for Nuclear Research (CERN) as a contractor. Frustrated with how his work was going, Berners-Lee essentially created the Web out of existing protocols and a working technology he had made earlier in order to make his job easier.

Berners-Lee had built a hypertext program in 1980 to help keep track of multiple projects and documentation. The system, called ENQUIRE, used hypertext to link project "cards" (we'd call them Web pages) to one another. He left CERN the same year, but returned in 1984 and played around with the program and the concept of hypertext until finally drafting a formal proposal for marrying the concept of hypertext to the Internet with the help of Robert Cailliau.

Berners-Lee pictured being able to navigate from document to document and across to different domains in the search of information. More importantly, he saw more than just textual information – the World Wide Web he envisioned would include colors and pictures, rather than just text. To do this, he needed to create the language of the documents (HTML), a system for addressing them (URL) and the protocols to reach them (HTTP).

URLs - www.yoursite.com

The Internet already had a system that Berners-Lee could build upon for his universal resource locator (URL). The domain name system (DNS) allowed computers to be reached by an easy-to-remember name rather than by a numerical IP address. In creating URLs, Berners-Lee used the domain names of the computers that hosted the hypertext files as the core address, and then used hierarchical slashes to separate the domain from specific files and folders.

So, someone looking for Berners-Lee's original Web page can find it on the World Wide Web Consortium's Web server by putting in the domain name (www.w3.org) and using slashes to get to the history folder (www.w3.org/History). From there, you can click on a graphical representation of a folder (199221103-hypertext/) and, by clicking more links, reachhttp://www.w3.org/History/19921103-hypertext/hypertext/WWW/TheProject.html - a reproduction of the original Web page.

Berners-Lee did not invent the idea of a link. Credit for that goes to Vannevar Bush's Memex, Ted Nelson's work in the '60s or Douglas Engelbart's initial demonstration of the concept. Berners-Lee did, however, spread the concept further than other inventors at the time. As much as we think of it now as a basic concept, it truly is one of the most important innovations in modern technology. Prior to the ability to link, there was no way to join together information from different parts of the world. (Learn more about Bush, Engelbart and Nelson in The Pioneers of the World Wide Web.)

HTTP

HyperText Transfer Protocol (HTTP) is the language that allows your machine to talk with other computers on the Internet. HTTP is a high-level protocol, which means that it delivers whole Web pages to your computer's browser. TCP/IP, on the other hand, delivers packets comprising the page's contents. When you open a Web page and enter a URL, you are seeing the delivery that HTTP has provided in the form of the Web page that appears in your browser.

When you click a link in a browser, you are actually making an HTTP request to a Web server. It then replies in the form of an HTTP response – usually containing the content of a Web page. The design of HTTP is based on this request and response concept and, although this is a simple process, therein lies its beauty. HTTP worked on every network that followed the TCP/IP protocol, regardless of the type of machine being used to accesses it. In other words, a Windows box could access information on a Unix box. As with linking, we think of this as a fairly basic concept, but this is actually a huge breakthrough in the history of computing.

HTML

HyperText Markup Language (HTML) wasn't a new idea anymore than a URL was. To his credit, Berners-Lee has always described his role as the father of the World Wide Web as one of pulling together what was already out there. Hypertext was being used in proprietary software by 1990, but it wasn't being used to navigate beyond a single program. Berners-Lee designed HTML usingStandard Generalized Markup Language (SGML) for the structural formatting, but with the addition of an anchor to link text to another document. If SGML was the smart, big sister, HTML was the much simpler, smaller brother.

With HTTP to find URLs and load up documents written in HTML, Berners-Lee had a complete framework for his web of resources, which could be interlinked and navigated using hypertext. Berners-Lee used one of Steve Jobs' Unix-based NeXT computers and turned it into a Web server, where he hosted his documents in HTML format. He also created a combination Web browser and editor (confusingly called WorldWideWeb) for people to access it.

The World Wide Web went live in 1991, and continued to be refined over time. Berners-Lee and CERN also made the decision on April 30, 1993, to make the World Wide Web free to anyone and everyone.

The Browser and the Modern Web

One of the primary reasons that the early iterations of the Internet weren't very attractive to people is that they required users to spend a lot of time scanning directories looking for files. Once a connection was established, the user could navigate menus and open files, but it was just text upon text upon text. There were some early programs aimed at helping index all the

available information, but the data was so vast that an index didn't help navigate to the file that was needed. Enter the Web browser.

Mosaic

The World Wide Web, like much of the Internet, was initially a playground for intellectuals. Marc Andreessen and Eric Bina changed that in a very big way. The two were working for the National Center for Supercomputing Applications (NSCA), creating a Web browser. There was no shortage of Web browsers at the time, but almost all of them were Unix-based and naturally text heavy. There was really no unified approach to things like images or sounds, and most browsers handled them by loading separate windows. Andreessen and Bina created Mosaic, a browser that broke this mold.

The Graphical Web

Mosaic was clearly a unique browser in that it worked across all platforms. This meant that it worked with the Windows computers that represented the vast majority of personal computers in the world. Second, it was built to display the World Wide Web as an integrated experience with text, images and sound, all activated through the graphical user interface (GUI). The

development of Mosaic influenced the development of HTML by defining things like the   tag for embedded images before any consensus was reached.

This somewhat heavy-handed approach had a positive effect in that the pages making up the WWW were forced to meet the capabilities of this new browser. Within the same window, a user could now see a photo of a bird next to text describing its species and then play a file of its song. Like email, this was a killer app that would sell.

Mosaic spread rapidly and prompted non-academics to start producing websites for the new Web surfers using Mosaic to access the Web. These new sites were differentiated by their URLs; government sites used a ".gov" domain name, academics and nonprofit organizations used ".org," and the commercial domains – which would soon come to dominate the Internet – used ".com". The number of websites grew from less than 1,000 to more 10,000 in 1994 alone, and continued to grow at a breakneck pace.

A Little Company called Netscape

Mosaic worked just fine, but Andreessen saw more potential. In May of 1994, he struck out and founded Netscape along with Jim Clark, a tech entrepreneur. As a private company, Netscape had the resources to move quickly, allowing it to release the first commercial Web browser – Mozilla 1.0 – by December of that year. The product was renamed Netscape Navigator in

subsequent versions.

To say Netscape was a success is a huge understatement. It was one of the fastest-growing companies in U.S. history. It went public in 1995, making the founders instant millionaires. Netscape advertised that "the Web is for everyone." It also represented the first time that average computer users could get online with relative ease.

It didn't take long for Microsoft to notice this new threat, and it fought back with a browser of its own by the name of [Internet Explorer](). By giving away the browser for free as part of the Windows operating system, Microsoft was able to undercut Netscape. By the end of 1998, Netscape was purchased by America Online (AOL), which eventually stopped putting resources into the company. So while Netscape is no longer around, its importance cannot be understated in the development of the modern Web.

When We Say Internet, We Mean the Web

When most of us talk about the Internet, we are really referring to the World Wide Web. The Internet is a system of protocols that allows computers to network, while the World Wide Web is what we are using that network to access. From the mid-1990s through the dot-com bubble, the Web really came into its own, eventually overshadowing the concept of an Internet entirely. However, without the groundwork done to create the Internet itself, there could be no Web.

The history of the Internet is a story of small achievements building up to something grand, something that puts the user at the center. Because of CERN's decision not to make the Web proprietary, we have the freedom to use it for educational purpose, just as the NSF intended, or to keep up on celebrity gossip (like celebrity gossip magazine TMZ intended). Whether it's renovation tips, the classics available via Project Gutenberg, or even the seedier underbelly that includes cyber crime and violent, discriminatory or pornographic websites, the Internet brings it all to you at the click of a link.

Vannevar Bush and others believed such a resource would make us much smarter by allowing us to store and access great swaths of knowledge outside ourselves, and generally leave mankind infinitely wiser and better off. Sometimes it feels like all that external storage has just made room for committing more sports trivia and quotes from "The Simpsons" to memory. Still, it is easy to see how we are better off thanks to the Internet, at least in terms of the free flow of information and communication in our daily lives. Many people who now spend hours a day online – whether for work or pleasure – have no idea how many tiny steps it took to make this a technology for the masses. (Learn more about the people who made the Internet possible in [The Pioneers of the World Wide Web]().)

The Future of the Internet

Now that we've had a bird's-eye view of all the work that went into it, the natural question is to ask what is in store for the future. We've already seen some of the future with social media, cloud computing and mobile devices that allow us to be connected to the Internet in all but the most remote locations. And there are small hints of what may or may not be some day. One of the most interesting questions is whether someone 20 years from now will even be reading this history, rather than listening to it being read by a browser or viewing it in a purely graphical representation of videos and infographics. Will the term browser still apply in 20 years? Will free content driven by advertising still be the primary business model?

These are all good questions, but ones that can only be answered by time. The one sure thing is that the Internet and the World Wide Web - and our relationship with the two - will continue to evolve in the future.