20 THINGS I LEARNED ABOUT BROWSERS AND THE WEB

Illustrated by Christoph Niemann. Written by the Google Chrome Team.
www.20thingsilearned.com

Foreword to 20 Things

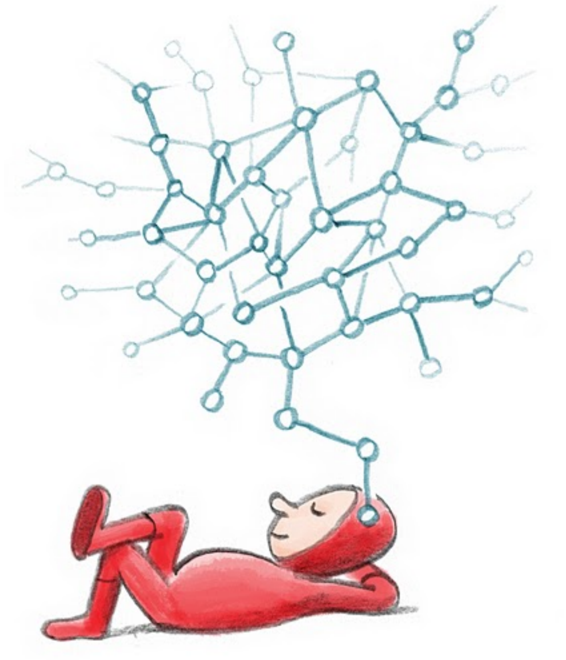

Many of us these days depend on the World Wide Web to bring the world's information to our fingertips, and put us in touch with people and events across the globe instantaneously. These powerful online experiences are possible thanks to an open web that can be accessed by anyone through a web browser, on any Internet-connected device in the world.

But how do our browsers and the web actually work? How has the World Wide Web evolved into what we know and love today? And what do we need to know to navigate the web safely and efficiently?

"20 Things I Learned About Browsers and the Web" is a short guide for anyone who's curious about the basics of browsers and the web. Here's what you'll find here:

First we'll look at the Internet, the very backbone that allows the web to exist. We'll also take a look at how the web is used today, through cloud computing and web apps.

Then, we'll introduce the building blocks of web pages like HTML and JavaScript, and review how their invention and evolution have changed the websites you visit every day. We'll also take a look at the modern browser and how it helps users browse the web more safely and securely.

Finally, we'll look ahead to the exciting innovations in browsers and web technologies that we believe will give us all even faster and more immersive online experiences in the future.

Life as citizens of the web can be liberating and empowering, but also deserves some self-education. Just as we'd want to know various basic facts as citizens of our physical neighborhoods -- water safety, key services, local businesses -- it's increasingly important to understand a similar set of information about our online lives. That's the spirit in which we wrote this guide. Many of the examples used to illustrate the features and functionality of the browser often refer back to Chrome, the open-source browser that we know well. We hope you find this guide as enjoyable to read as we did to create.

Happy browsing!

The Google Chrome Team, with many thanks to Christoph Niemann for his illustrations
November 2010

Thing #1:

What is the Internet?

or, “You Say Tomato, I Say TCP/IP”

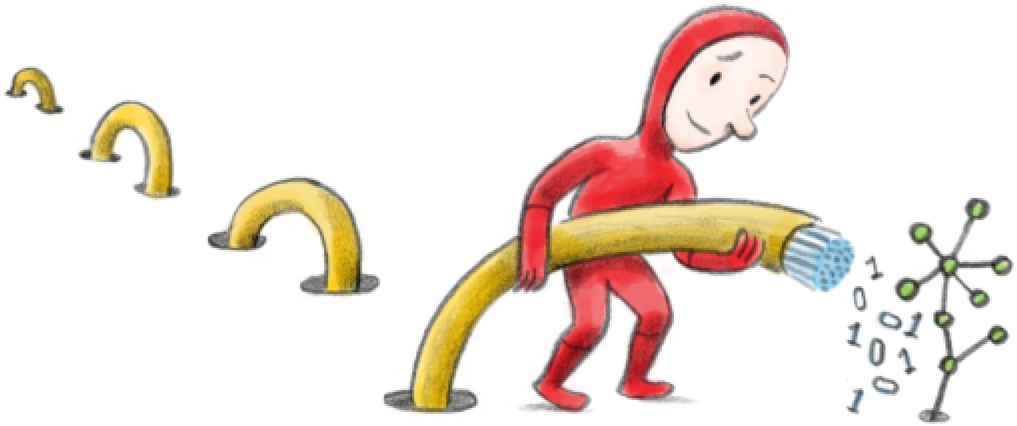

What is the Internet, exactly?

To some of us, the Internet is where we stay in touch with friends, get the news, shop, and play games. To some others, the Internet can mean their local broadband providers, or the underground wires and fiber-optic cables that carry data back and forth across cities and oceans. Who is right?

A helpful place to start is near the Very Beginning: 1974. That was the year that a few smart computer researchers invented something called the Internet Protocol Suite, or TCP/IP for short. TCP/IP created a set of rules that allowed computers to “talk” to each other and send information back and forth.

TCP/IP is somewhat like human communication: when we speak to each other, the rules of grammar provide structure to language and ensures that we can understand each other and exchange ideas. Similarly, TCP/IP provides the rules of communication that ensure interconnected devices understand each other so that they can send information back and forth. As that group of interconnected devices grew from one room to many rooms -- and then to many buildings, and then to many cities and countries -- the Internet was born.

The early creators of the Internet discovered that data and information could be sent more efficiently when broken into smaller chunks, sent separately, and reassembled. Those chunks are called **packets**. So when you send an email across the Internet, your full email message is broken down into packets, sent to your recipient, and reassembled. The same thing happens when you watch a video on a website like YouTube: the video files are segmented into data packets that can be sent from multiple YouTube servers around the world and reassembled to form the video that you watch through your browser.

What about speed? If traffic on the Internet were akin to a stream of water, the Internet's **bandwidth** is equivalent to the amount of water that flows through the stream per second. So when you hear engineers talking about bandwidth, what they're really referring to is the amount of data that can be sent over your Internet connection per second. This is an indication of how fast your connection is. Faster connections are now possible with better physical infrastructure (such as fiber optic cables that can send information close to the speed of light), as well as better ways to encode the information onto the physical medium itself, even on older medium like copper wires.

The Internet is a fascinating and highly technical system, and yet for most of us today, it's a user-friendly world where we don't even think about the wires and equations involved. The Internet is also the backbone that allows the World Wide Web that we know and love to exist: with an Internet connection, we can access an open, ever-growing universe of interlinked web pages and applications. In fact, there are probably as many pages on the web today as there are neurons in your brain, as there are stars in the Milky Way!

In the next two chapters, we'll take a look at how the web is used today through cloud computing and web apps.

Thing #2:

Cloud Computing

or, why it's ok for a truck to crush your laptop

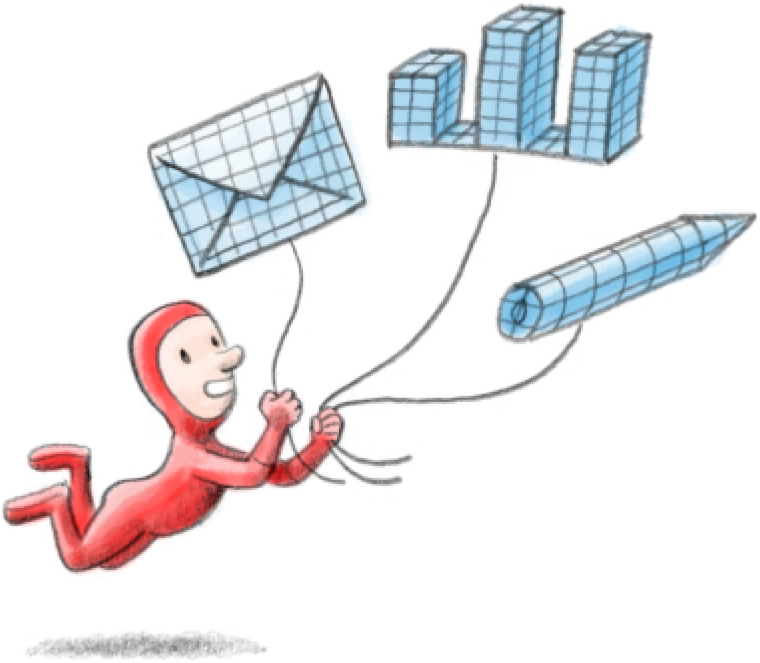

Modern computing in the age of the Internet is quite a strange, remarkable thing. As you sit hunched over your laptop at home watching a YouTube video or using a search engine, you're actually plugging into the collective power of thousands of computers that serve all this information to you from far-away rooms distributed around the world. It's almost like having a massive supercomputer at your beck and call, thanks to the Internet.

This phenomenon is what we typically refer to as cloud computing. We now read the news, listen to music, shop, watch TV shows and store our files on the web. Some of us live in cities in which nearly every museum, bank, and government office has a website. The end result? We spend less time in lines or on the phone, as these websites allow us to do things like pay bills and make reservations. The movement of many of our daily tasks online enables us to live more fully in the real world.

Cloud computing offers other benefits as well. Not too long ago, many of us worried about losing our documents, photos and files if something bad happened to our computers, like a virus or a hardware malfunction. Today, our data is migrating beyond the boundaries of our personal computers. Instead, we're moving our data online into "the cloud". If you upload your photos, store critical files online and use a web-based email service like Gmail or Yahoo! Mail, an 18-wheel truck could run over your laptop and all your data would still safely reside on the web, accessible from any Internet-connected computer, anywhere in the world.

Thing #3:

Web Apps

or, “Life, Liberty and the Pursuit of Appiness”

If you play online games, use an online photo editor, or rely on web-based services like Google Maps, Twitter, Amazon, YouTube or Facebook, then you’re an active resident in the wonderful world of web apps.

What exactly *is* a web app, anyway? And why should we care?

App is shorthand for an application. Applications are also called programs or software. Traditionally, they’ve been designed to do broad, intensive tasks like accounting or word processing. In the online world of web browsers and smart phones, apps are usually nimbler programs focused on a single task.

Web apps, in particular, run these tasks inside the web browser and often provide a rich, interactive experience.

Google Maps is a good example of a web app. It's focused on one task: providing helpful map features within a web browser. You can pan and zoom around a map, search for a college or cafe, and get driving directions, among other tasks. All the information you need is pulled into the web app dynamically every time you ask for it.

This brings us to four virtues of Web Appiness:

1. I can access my data from anywhere.

In the traditional world of desktop applications, data is usually stored on my computer's hard drive. If I'm on vacation and leave my computer at home, I can't access my email, photos, or any of my data when I need it. In the new world of web apps, my email and all my data are stored online on the web. I can get to it on a web browser from any computer that's connected to the Internet.

2. I'll always get the latest version of any app.

Which version of YouTube am I using today? What about tomorrow? The answer: Always the latest. Web apps update themselves automatically, so there's always just one version: the latest version, with all the newest features and improvements. No need to manually upgrade to a new version every time. And I don't have to go through a lengthy install process to use my web apps.

3. It works on every device with a web browser.

In traditional computing, some programs work only on particular systems or devices. For instance, many programs written for a PC won't work on a Mac. Keeping up with all the right software can be time-consuming and costly. In contrast, the web is an open platform. Anyone can reach it from a browser on any web-connected device, regardless of whether it's a desktop computer, laptop, or mobile phone. That means I can use my favorite web apps even if I'm using my friend's laptop or a computer at an Internet cafe.

4. It's safer.

Web apps run in the browser and I never have to download them onto my computer. Because of this separation between the app code and my computer's code, web apps can't interfere with other tasks on my computer or the overall performance of my machine. This means that I'm better protected from threats like viruses, malware and spyware.

Thing #4:

HTML, JavaScript, CSS and more

or, This is not your mom’s AJAX

Web pages are written in HTML, the web programming language that tells browsers how to lay out and present content on a web page. In other words, HTML provides the basic building blocks for the web. For a long time, those building blocks were pretty simple and static: lines of text, links and images.

Today, the web goes beyond just text, links, and images. We expect to be able to play online chess or seamlessly scroll around a map of our neighborhood, without waiting for the entire page to reload for every chess move or every map scroll.

The idea of such dynamic web pages began with the invention of the scripting language JavaScript. JavaScript support in major web browsers meant that web pages could incorporate more meaningful real-time interactions. For example, if you’ve filled out an online form and hit the “submit” button, the web page can use JavaScript to check your entries in real-time and alert you almost instantly if you had filled out the form wrongly.

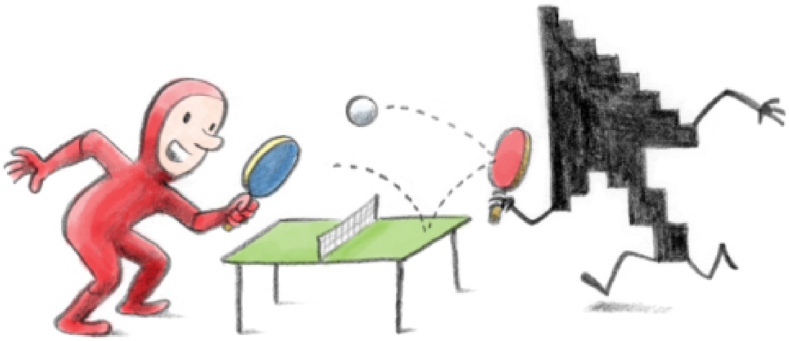

But the dynamic web as we know it today truly came to life when XHR (XMLHttpRequest) was introduced into JavaScript, and first used in web applications like Microsoft Outlook for the Web, Gmail and Google Maps. XHR enabled individual parts of a web page -- a game, a map, a video, a little survey -- to be altered without needing to reload the entire page. As a result, web apps are faster and more responsive.

Web pages have also become more expressive with the introduction of CSS (Cascading Style Sheets). CSS gives programmers an easy, efficient way to define a web page's layout and beautify the page with design elements like colors, rounded corners, gradients, and animation.

Web programmers often refer to this potent combination of JavaScript, XHR, CSS and several other web technologies as AJAX (Asynchronous JavaScript and XML). HTML has also continued to evolve as more features and improvements are incorporated into new versions of the HTML standard.

Today’s web has evolved from the ongoing efforts of all the technologists, thinkers, coders and organizations who create these web technologies and ensure that they’re supported in web browsers like Internet Explorer, Firefox, Safari and Google Chrome. This interaction between web technologies and browsers has made the web an open and friendly construction platform for web developers, who then bring to life many useful and fun web applications that we use daily.

Thing #5:

HTML5

or, in the beginning, there was no <video>

More than two decades after HTML was introduced, we're still asking questions about what the web is, and what it might become. What kinds of features and applications would we, as users, find fun, useful or even indispensable? What tools do developers need in order to create these great sites and apps? And finally, how can all this goodness be delivered inside a web browser?

These questions led to the evolution of the latest version of HTML known as HTML5, a set of capabilities that gives web designers and developers the ability to create the next generation of great online applications. Take the HTML5 <video> tag, for example. Video wasn't a major (or, really, any) part of the early web; instead, internet users installed additional software called plug-ins, in order to watch videos inside their web browsers. Soon it became apparent that easy access to video was a much-wanted feature on the web. The introduction of the <video> tag in HTML5 allows videos to be easily embedded and played in web pages without additional software.

Other cool HTML5 features include offline capabilities that let users interact with web apps even when they don't have an internet connection, as well as drag-and-drop capabilities. In Gmail, for instance, easy drag-and-drop allows users to instantly attach a file to an email message by simply dragging the file from the user's desktop computer into the browser window.

HTML5, like the web itself, is in perpetual evolution, based on users' needs and developers' imaginations. As an open standard, HTML5 embodies some of the best aspects of the web: it works everywhere, and on any device with a modern browser. But just as you can only watch HDTV broadcasts on an HD-compatible television, you need to use an up-to-date, HTML5-compatible browser in order to enjoy sites and apps that take advantage of HTML5's features. Thankfully, as an Internet user, you have lots of choice when it comes to web browsers -- and unlike TVs, web browsers can be downloaded for free.

Thing #6:

3D in the Browser

or, browsing with more depth

3D graphics and animation can be truly captivating with all the right details in place: details like lighting and shadows, reflections, and realistic textures. But until now, it has been hard to deliver a compelling 3D experience, particularly over the Internet.

Why? Mostly because creating a 3D experience in games and other applications requires data -- lots and lots of data -- to display intricate textures and shapes. In the past, these large amounts of data demanded more Internet bandwidth and more computing power than most common systems could handle.

All that has changed very recently, and all for the better: browser-based 3D has arrived.

Modern broadband helped solve bandwidth needs. Many homes and offices now have broadband speeds that dwarf the connections of even ten years ago. As a result, it's possible to send large amounts of data over the Internet -- data that is needed to display realistic 3D experiences in the browser. In addition, the computers we use today are so much more powerful than what we had in the past: processors and memory have improved such that even a standard laptop or desktop today can handle the complexity of 3D graphics.

Neither broadband nor raw computing power would matter without substantial advancements in the web browser's capabilities. Many modern browsers have adopted open web technologies like WebGL and 3D CSS. With these technologies, web developers can create cool 3D effects for their web applications, and we can experience them without needing additional plug-ins. On top of that, many modern browsers now take advantage of a technique known as hardware-acceleration. This means that the browser can use the Graphics Processing Unit, or GPU, to speed up the computations needed to display both 3D and everyday 2D web content.

So why is 3D in the browser a big deal? Because now it joins HTML5, JavaScript and other nifty new technologies in the toolkit that web developers can use to create a powerful new generation of web applications. For users, this means great new ways to visualize the information we find useful, and more fun online with engaging 3D environments and games.

Most importantly, 3D in the browser comes with all the goodness of web apps: you can share, collaborate, and personalize the latest apps with friends all over the world. Definitely more data and fun that *everyone* can use.

Thing #7:

A Browser Madrigal

or, old vs. modern browsers

Crabbed old and modern browsers
Cannot live together:
The modern browser is faster, featureful, and more secure
The old browser is slow, and at worst, a dreadful danger
Malicious attacks it cannot endure.

(with apologies to Shakespeare)

Most of us don't realize how much an old and out-of-date web browser can negatively impact our online lives, particularly our online safety. You wouldn't drive an old car with bald tires, bad brakes, and an unreliable engine for years on end. It's a bad idea to take the same chances with the web browser that you use daily to navigate to every page and application on the web.

Upgrading to a modern browser -- like the latest version of Mozilla Firefox, Apple Safari, Microsoft Internet Explorer, Opera, or Google Chrome -- is important for three reasons:

First, old browsers are vulnerable to attacks, because they typically aren't updated with the latest security fixes and features. Browser vulnerabilities can lead to stolen passwords, malicious software snuck secretly onto your computer, or worse. An up-to-date browser helps guard against security threats like phishing and malware.

Second, the web evolves quickly. Many of the latest features on today's websites and web applications won't work with old browsers. Only up-to-date browsers have the speed improvements that let you run web pages and applications quickly, along with support for modern web technologies such as HTML5, CSS3, and fast JavaScript.

Third and last, old browsers slow down innovation on the web. If lots of Internet users cling to old browsers, web developers are forced to design websites that work with both old and new technologies. Facing limited time and resources, they end up developing for the lowest common denominator -- and not building the next generation of useful, groundbreaking web applications. (Imagine if today's highway engineers were required to design high-speed freeways that would still be perfectly safe for a Model T.) That's why outdated browsers are bad for users overall and bad for innovation on the web.

Not that anyone blames you personally for staying loyal to your aging browser. In some cases, you may be unable to upgrade your browser. If you find that you're blocked from upgrading your browser on your corporate computer, have a chat with your IT administrator. If you can't upgrade an old version of Internet Explorer, the Google Chrome Frame plug-in can give you the benefits of some modern web app functionality by bringing in Google Chrome's capabilities into Internet Explorer.

Old, outdated browsers are bad for us as users, and they hold back innovation all over the web. So take a moment to make sure that you've upgraded to the latest version of your favorite modern browser.

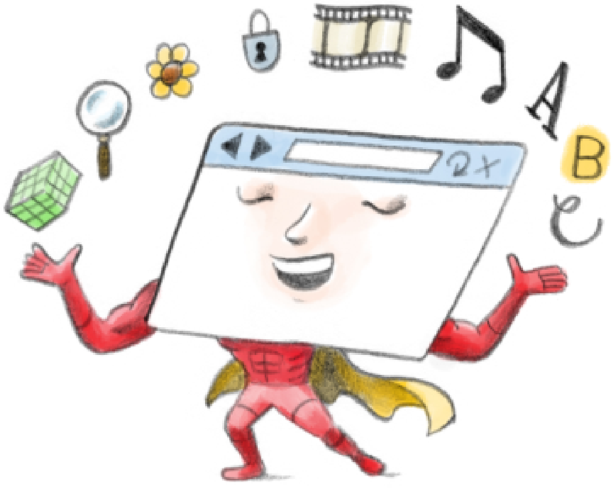

Editor's note: At the time of publication, the latest stable versions of the major modern browsers are

Firefox 3.6, Safari 5, Google Chrome 7, Internet Explorer 8, and Opera 10.63. To check which browser you're using, visit www.whatbrowser.org.

Thing #8:

Plug-ins

or, pepperoni for your cheese pizza

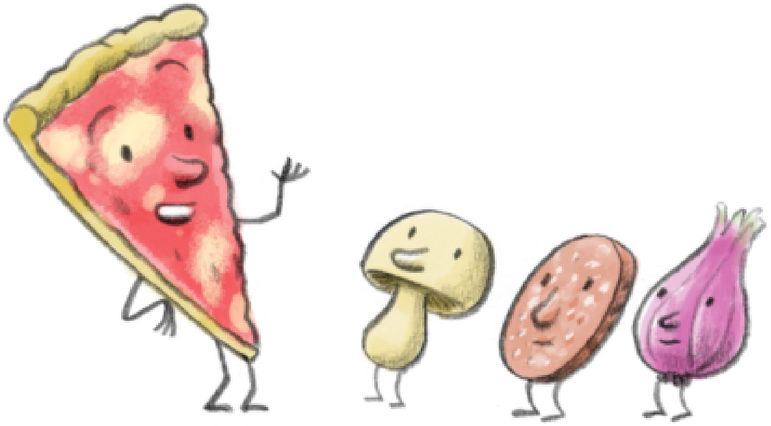

In the early days of the World Wide Web, the first versions of HTML couldn't deliver fancy content like videos. Text, images, and links were pretty much the limit.

Plug-ins were invented to work around the limitations of early HTML and deliver more interactive content. A plug-in is an additional piece of software that specializes in processing particular types of content. For example, users may download and install a plug-in like Adobe Flash Player to view a web page which contains a video or an interactive game.

How much does a plug-in interface with a browser? Curiously, hardly at all. The plug-in model is a lot like picture-in-a-picture on TV: the browser defines a distinct space on the web page for the plug-in, then steps aside. The plug-in is free to operate inside that space, independent of the browser.

This independence means that a particular plug-in can work across many different browsers. However, that ubiquity also makes plug-ins prime targets for browser security attacks. Your computer is even more vulnerable to security attacks if you're running plug-ins that aren't up to date, because out-of-date plug-ins don't contain the latest security fixes.

The plug-in model we use today is largely the one inherited from the web's early days. But the web community is now looking at new ways to modernize plug-ins -- like clever ways to integrate plug-ins more seamlessly so that their content is searchable, linkable, and can interact with the rest of the web page. More importantly, some browser vendors and plug-in providers now collaborate to protect users from security risks. For example, the Google Chrome and Adobe Flash Player teams have worked together to integrate Flash Player into the browser. Chrome's auto-update mechanism helps ensure that the Flash Player plug-in is never out-of-date and always receives the latest security fixes and patches.

Thing #9:

Browser Extensions

or, superpowers for your browser

Browser extensions let you add new features to your browser -- literally *extending* your browser.

This means that you can customize your browser with the features that are most important to you. Think of extensions as ways of adding new superpowers to what the browser can already do.

These superpowers can be mighty or modest, depending on your needs. For example, you might install a currency converter extension that shows up as a new button next to your browser's address bar. Click the button and it converts all the prices on your current web page into any currency you specify. That's helpful if you're an avid international backpacker who does most of your travel planning and booking online. Extensions like these let you apply the same kind of functionality to every web page you visit.

Browser extensions can also act on their own, outside of web pages. An email notifier extension can live on your browser toolbar, quietly check for new messages in your email account and let you know when one arrives. In this case, the extension is always working in the background no matter what web page you're looking at -- and you don't have to log in to your email in a separate window to see if you have new messages.

When browser extensions were first introduced, developers often had to build them in unusual programming languages or in heavy-duty mainstream languages like C++. This took a lot of work, time and expertise. Adding more code to the browser also added to security concerns, as it gave attackers more chances to exploit the browser. Because the code was sometimes arcane, extensions were notorious for causing browser crashes, too.

Today, most browsers let developers write extensions in the basic, friendly programming languages of the web: HTML, JavaScript and CSS. Those are the same languages used to build most modern web apps and web pages, so today's extensions are much closer cousins to the web apps and pages they work with. They're faster and easier to build, safer, and get better and better right along with the web standards they're built upon.

To discover new extensions, check out your browser's extensions gallery. You'll see thousands of extensions that can help make browsing more efficient or just plain fun -- from extensions that let you highlight and scribble notes on web pages while you're doing research, to those that show nail-biting, play-by-play sports updates from your browser's interface.

Thing #10:

Synchronizing the Browser

or, why it's ok for a truck to crush your laptop, part II

So you're living in "the cloud": congratulations! You use web apps for email, music, and almost everything. You save critical documents, photos, and files online where you can reach them from any Internet-connected computer, anywhere in the world.

If an 18-wheel truck comes roaring down the road and crushes your laptop to bits, all is not lost. You just find another Internet-connected device and get back to working with all that vital information you so smartly saved online.

But wait: What about all the bookmarks, browser extensions, and browser preferences that you use daily? Did they get crunched into oblivion along with your laptop?

The answer used to be "yes." You'd have to forage for your favorite extensions all over again and gather all the websites you had painstakingly bookmarked. But no more! Many of today's browsers, such as Firefox and Chrome, have begun building in a feature known as **synchronization** ("sync" for short). Sync lets you save your browser settings online, in the cloud, so they aren't lost even if your computer melts down.

Sync functionality also makes life simpler if you use multiple computers, say, a laptop at work and a family desktop at home. You don't have to manually recreate bookmarks of your favorite websites or reconfigure the browser settings on every computer you own. Any changes you make to your sync-enabled browser on one computer will automatically appear in all other synced computers within seconds.

In Chrome, for example, sync saves all bookmarks, extensions, preferences and themes to your Google Account. Use any other Internet-connected computer, and all you need to do is fire up Chrome and log in to your Google Account through the browser’s sync feature. *Voila!* All your favorite browser settings are ready to use on the new machine.

Regardless of how many computers you need to juggle, as long as you have an Internet connection and a modern browser that’s synced to the cloud, you’re all set to go. Even if every one of them gets hit by the proverbial truck.

Thing #11:

Browser Cookies

or, thanks for the memories

Cookie seems like an unlikely name for a piece of technology, but cookies play a key role in providing functionality that Internet users may want from websites: a memory of visits, in the past or in progress.

A cookie is a small piece of text sent to your browser by a website you visit. It contains information about your visit that you may want the site to remember, like your preferred language and other settings. The browser stores this data and pulls it out the next time you visit the site to make the next trip easier and more personalized. If you visit a movie website and indicate that you're most interested in comedies, for instance, the cookies sent by the website can remember this so you may see comedies displayed at the start of your next visit.

Online shopping carts also use cookies. As you browse for DVDs on that movie shopping site, for instance, you may notice that you can add them to your shopping cart without logging in. Your shopping cart doesn't "forget" the DVDs, even as you hop around from page to page on the shopping site, because they're preserved through browser cookies. Cookies can be used in online advertising as well, to remember your interests and show you related ads as you surf the web.

Some people prefer not to allow cookies, which is why most modern browsers give you the ability to manage cookies to suit your tastes. You can set up rules to manage cookies on a site-by-site basis, giving you greater control over your privacy. What this means is that you can choose which sites you trust and allow cookies only for those sites, blocking cookies from everyone else. Since there are many types of cookies -- including "session-only cookies" that last only for a particular browsing session, or permanent cookies that last for multiple sessions -- modern browsers typically give you fine-tuned controls so that you can specify your preferences for different types of cookies, such as accepting permanent cookies as session-only.

In the Google Chrome browser, you’ll notice a little something extra in the Options menus: a direct link to the Adobe Flash Player storage settings manager. This link makes it easy to control local data stored by Adobe Flash Player (sometimes referred to as "Flash cookies"), which can contain information on Flash-based websites and applications that you visit. Just as you can manage your browser cookies, you should be able to easily control your Flash cookies settings as well.

Thing #12:

Browsers and Privacy

or, giving you choices to protect your privacy in the browser

Security and privacy are closely related, but not identical.

Consider the security and privacy of your home: door locks and alarms help protect you from burglars, but curtains and blinds keep your home life private from passersby. In the same way, browser security helps protect you from malware, phishing, and other online attacks, while privacy features help keep your browsing private on your computer.

Let’s look more closely at privacy. Here’s an analogy: Say you’re an avid runner who jogs a few miles every day. If you carry a GPS device to help you track your daily runs, you create a diary of running data on your device -- a historical record of where you run, how far you run, your average speed, and the calories you burn.

As you browse the web, you generate a similar diary of *browser* data that is stored locally on your computer: a history of the sites you visit, the cookies sent to your browser, and any files you download. If you’ve asked your browser to remember your passwords or form data, that’s stored on your computer too.

Some of us may not realize that we can clear all this browser data from our computers at any time. It’s easy to do through a browser’s Options or Preferences menu. (The menu differs from browser to browser.) In fact, the latest versions of most modern browsers also offer a “private” or “incognito” mode. For example, in Chrome’s incognito mode, any web page that you view won’t appear in your browsing history. In addition, all new cookies are deleted after you close all the incognito windows that you’ve opened. This mode is especially handy if you share your computer with other people, or if you work on a public computer in your local library or cybercafe.

All these privacy features in the browser give you control over the browsing data locally on your computer or specific data that are sent by your browser to websites. Your browser's privacy settings do not control other data that these websites may have about you, such as information you previously submitted on the website.

There are ways to limit some of the information that websites receive when you visit them. Many browsers let you control your privacy preferences on a site-by-site basis and make your own choices about specific data such as cookies, JavaScript, and plugins. For instance, you can set up rules to allow cookies only for a specified list of sites that you trust, and instruct the browser to block cookies for all other sites.

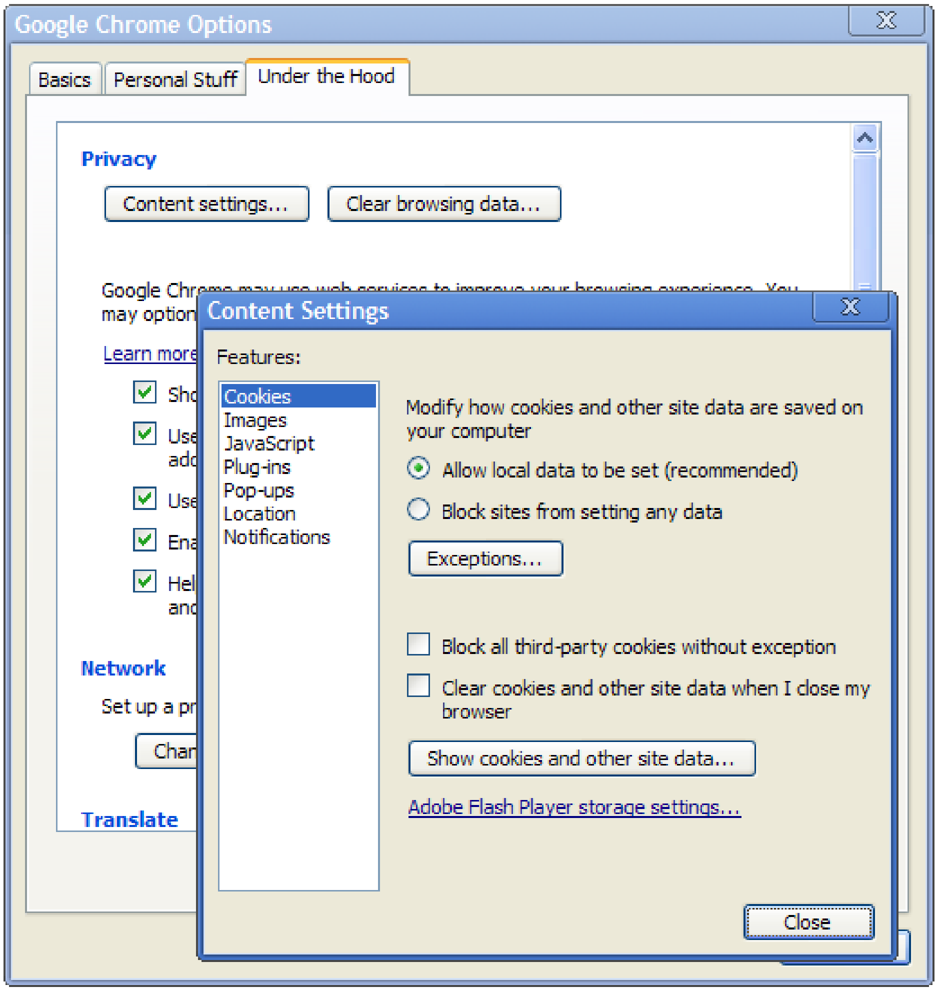

Modify how cookies and other site data are saved on your computer

● Allow local data to be set (recommended)

○ Block sites from setting any data

□ Block all third-party cookies without exception

□ Clear cookies and other site data when I close my browser

Example of privacy controls in the browser

There's always a bit of tension between privacy and efficiency. Collecting real-world aggregate data and feedback from users can really help improve products and the user experience. The key is finding a good balance between the two while upholding strong privacy standards.

Here's an example from the real world: browser cookies. On one hand, with cookies, a website you frequently visit is able to remember contents of your shopping cart, keep you logged in, and deliver a more useful, personalized experience based on your previous visits. On the other hand, allowing browser cookies means that the website is collecting and remembering information about these previous visits. If you wish, you can choose to block cookies at any time. So the next time you’re curious about fine-tuning your browser privacy settings, check out the privacy settings in your browser’s Options or Preferences menu.

Thing #13:

Malware, Phishing, and Security Risks

or, “If it quacks like a duck but isn’t a duck”

When you use an ATM downtown, you probably glance over your shoulder to make sure nobody is lurking around to steal your PIN number (or your cash). In fact, you probably first check to make sure that you’re not using a fake ATM machine. When you browse the web and perform transactions online, two security risks to be aware of are malware and phishing. These attacks are perpetrated by individuals or organizations who hope to steal your personal information or hijack your computer.

What exactly are phishing and malware attacks?

Phishing takes place when someone masquerades as someone else, often with a fake website, to trick you into sharing personal information. (It’s called “phishing” because the bad guys throw out electronic bait and wait for someone to bite.) In a typical phishing scam, the attacker sends an email that looks like it’s from a bank or familiar web service you use. The subject line might say, “Please update your information at your bank!” The email contains phishing links that look like they go to your bank’s website, but really take you to an impostor website. There you’re asked to log in, and inadvertently reveal your bank account number, credit card numbers, passwords, or other sensitive information to the bad guys.

Malware, on the other hand, is malicious software installed on your machine, usually without your knowledge. You may be asked to download an anti-virus software that is actually a virus itself. Or you may visit a page that installs software on your computer without even asking. The software is really designed to steal credit card numbers or passwords from your computer, or in some cases, harm your computer. Once the malware is on your computer, it’s not only difficult to remove, but it’s also free to access all the data and files it finds, send that information elsewhere, and generally wreak havoc on your computer.

An up-to-date, modern web browser is the first line of defense against phishing and malware attacks. Most modern browsers, for instance, can help analyze web pages to look for signs of lurking malware, and alert you when they find it.

At the same time, an attacker may not always use sophisticated technical wizardry to hijack your computer, but could instead find clever ways to trick you into making a bad decision. In the next few chapters, we’ll look at how you can make wiser decisions to protect yourself when you’re online -- and how browsers and other web technologies can help.

Thing #14:

How Modern Browsers Help Protect You from Malware and Phishing

or, beware the ne’er-do-wells!

An up-to-date browser guards you from phishing and malware attacks when you’re browsing the web. It does so by limiting three types of security risk when you’re online:

Risk 1: How often you come into contact with an attacker

You can be exposed to attackers through a malicious fake website, or even through a familiar website that has been hacked. Most modern browsers pre-check each web page you visit and alert you if one is suspected of being malicious. This lets you make an informed judgment about whether you really want to visit that page.

For example, Google Chrome uses Safe Browsing technology, which is also used in several other modern browsers. As you browse the web, each page is checked quickly against a list of suspected phishing and malware websites. This list is stored and maintained locally on your computer to help protect your browsing privacy. If a match against the local list is found, the browser then sends a request to Google for more information. (This request is completely obscured and the browser does not send it in plain text.) If Google verifies the match, Chrome shows a red warning page to alert you that the page you're trying to visit may be dangerous.

Risk 2: How vulnerable your browser is if it's attacked

Old browsers that haven't been upgraded are likely to have security vulnerabilities that attackers can exploit. All outdated software, irrespective of whether it's your operating system, browser, or plug-ins, has the same problem. That's why it's important to use the very latest version of your browser and promptly install security patches on your operating system and all plug-ins, so that they're always up-to-date with the latest security fixes.

Some browsers check for updates automatically and install updates when initiated by the user. Chrome and some other browsers go one step further: they're built with auto-update. The browser runs an update check periodically, and automatically updates to the latest version without disrupting your browsing flow. Furthermore, Chrome has integrated Adobe Flash Player and a PDF viewer into the browser, so that both these popular plug-ins are also auto-updated.

Risk 3: How much damage is done if an attacker finds vulnerabilities in your browser

Some modern browsers like Chrome and Internet Explorer are built with an added layer of protection known as a "sandbox." Just as a real-life sandbox has walls to keep sand from spilling out, a browser sandbox builds a contained environment to keep malware and other security threats from infecting your computer. If you open a malicious web page, the browser's sandbox prevents that malicious code from leaving the browser and installing itself to your hard drive. The malicious code therefore cannot read, alter, or further damage the data on your computer.

In summary, a modern browser can protect you against online security threats by first, checking websites you're about to visit for suspected malware and phishing; second, providing update notifications or autoupdating when a newer, more secure version of the browser is available, and third, using the browser sandbox to curb malicious code from causing further damage to your computer.

In the next few chapters, we'll take a look at how a basic understanding of web addresses can help you make informed decisions about the websites you visit.

Thing #15:

Using Web Addresses to Stay Safe

(or, “My name is URL”)

A Uniform Resource Locator -- better known as a URL -- may sound like a complicated thing. But fret not: it’s simply the web address you type into your browser to get to a particular web page or web application.

When you enter a URL, the website is fetched from its hosting server somewhere in the world, transported over miles of cables to your local Internet connection, and finally displayed by the browser on your computer. Here are a few examples of a URL:

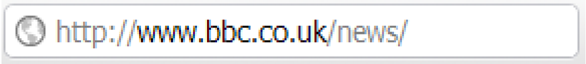
http://www.bbc.co.uk/news/

...to get to the news website for the British Broadcasting Corporation
(“.co.uk” indicates registration in the United Kingdom)

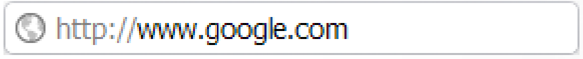
http://www.google.com

...to get to the search engine Google

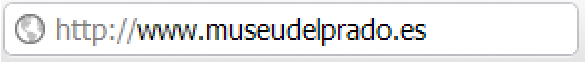
http://www.museudelprado.es

...to get to the website for Museo Nacional Del Prado, the Madrid-based art museum.
(“.es” indicates registration in Spain)

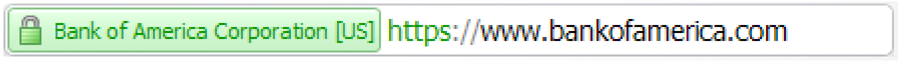

Bank of America Corporation [US] https://www.bankofamerica.com

...to get to the online banking website for Bank of America
(“https://” indicates an encrypted connection)

It’s easy to take URLs for granted, since we type them into our browsers every day. But understanding the parts of a URL can help guard against phishing scams or security attacks.

Let’s look at what’s in a URL in this example:

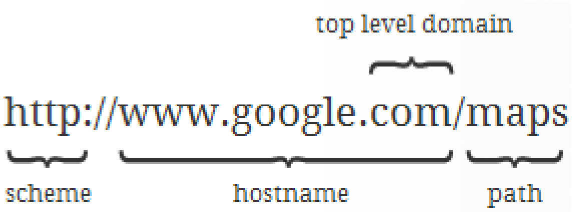

The first part of a URL is called the **scheme**. In the example above, HTTP is the scheme and shorthand for HyperText Transfer Protocol.

Next, “www.google.com” is the name of the **host** where the website resides. When any person or company creates a new web site, they register this hostname for themselves. *Only* they may use it. This is important, as we’ll see in a moment.

A URL may have an additional **path** after the hostname, which sends you to a specific page on that host -- like jumping right to a chapter or page in a book. Back to our example, the path tells the host server that you want to see the maps web application at www.google.com. (In other words, Google Maps.) Sometimes that path is moved to the front of the hostname as a full subdomain, such as “maps.google.com”, or “news.google.com” for Google News.

Now let’s talk safety. One way to check if you’re surfing right into a phishing scam or an impostor website is by looking carefully at the URL in your browser’s address bar. Pay particular attention to the hostname -- remember, only the legitimate owner of that hostname can use it.

For example, if you click on a link and expect to be directed to the Bank of America website:

LEGITIMATE:

- www.bankofamerica.com is a legitimate URL, since the hostname is correct.
- www.bankofamerica.com/smallbusiness is also a legitimate URL since the hostname is correct. The path of the URL points to a sub-page on small business.

SUSPICIOUS:

- bankofamerica.xyz.com is not Bank of America’s website. Instead, “bankofamerica” is a subdomain of the website xyz.com.
- www.xyz.com/bankofamerica is still not Bank of America's website. Instead, “bankofamerica” is a path within www.xyz.com.

If you’re using a banking website or conducting an online transaction with sensitive information such as your password or account number, check the address bar first! Make sure that the scheme is “**https:/**

/" and there's a padlock icon in your browser's address bar. "https://" indicates that the data is being transported between the server and browser using a secure connection.

Through a secure connection, the full URL for Bank of America's website should look like this: **https://** www.bankofamerica.com. A secure connection ensures that no one else is eavesdropping or interfering with the sensitive information that you're sending. So "https://" is a good sign. But remember, it's still important to make sure that you're actually talking to a legitimate website by checking the hostname of a URL. (It would defeat the purpose to have a secure connection to a bogus website!)

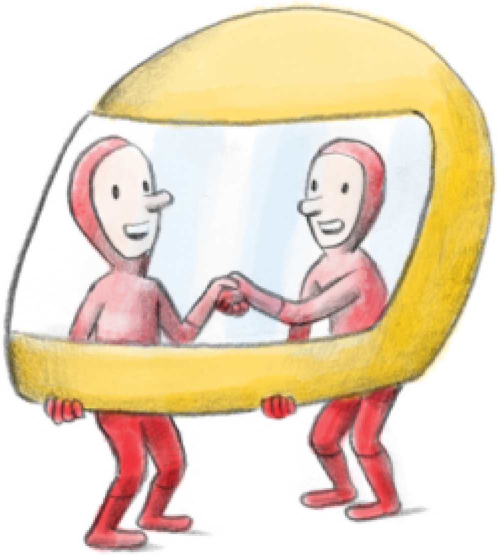

In the next chapter, we'll look at how a typed URL into the browser's address bar takes you to the right web page.

Thing #16:

IP Addresses and DNS

or, the phantom phone booth

Do you wonder how your browser finds the right web page when you type a URL into its address bar?

Every URL (say, "www.google.com") has its own numbered Internet Protocol or IP address. An IP address looks something like this:

74.125.19.147

An IP address is a series of numbers that tells us where a particular device is on the Internet network, be it the google.com server or your computer. It's a bit like mom's phone number: just as the phone number tells an operator which house to route a call to so it reaches your mom, an IP address tells your computer which other device on the Internet to communicate with -- to send data to and get data from.

Your browser doesn't automatically know every IP address for the 35 billion (or more) devices on the planet that are connected on the Internet. It has to look each one up, using something called the Domain Name System. The DNS is essentially the "phone book" of the Web: while a phone book translates a name like "Acme Pizza" into the right phone number to call, the DNS translates a URL or web address (like "www.google.com") into the right IP address to contact (like "74.125.19.147") in order to get the information that you want (in this case, the Google homepage).

So when you type in “google.com” into your web browser, the browser looks up google.com’s IP address through a DNS and contacts it, waits for a response to confirm the connection, and then sends your request for google.com’s web page to that IP address. Google’s server at that IP address will then send back the requested web page to your computer’s IP address for your browser to display.

In many ways, fetching and loading a web page in the browser is not unlike making a phone call. When you make a phone call, you’d probably look up the number, dial, wait for someone to pick up, say “hello,” and wait for a response before you start the conversation. Sometimes you have to redial if there are problems connecting. On the web, a similar process happens in a split-second; all you see is that you’ve typed “www.google.com” into the browser and the Google home page appears.

In the next chapter, we’ll look at how we can verify the identity of a website that we fetch and load in the browser through the **extended validation certificate**.

Thing #17:

Validating Identities Online

or, “Dr. Livingstone, I presume?”

In the physical world, you can see the people you share information with. You talk to them face-to-face, or meet them in a trusted place like a bank branch. That’s how you make your first judgments about giving them your trust.

But online, it can be hard to tell who’s behind any website. The visual cues we normally rely on can be faked. For example, a phony webpage could copy the logo, icon, and design of your own bank’s website -- almost as if they had set up a fake storefront on your block.

Fortunately, there are tools to help you determine if a website is genuine or not. Some websites have an **extended validation certificate** that allows you to determine the name of the organization that runs the web site. The extended validation certificate gives you the information you need to help ensure that you’re not entrusting your information to a fake website.

Here’s an example of extended validation in action in the browser. On a bank’s website that has been verified through extended validation, the bank’s name is displayed in a green box between the lock icon and the web address in the address bar:

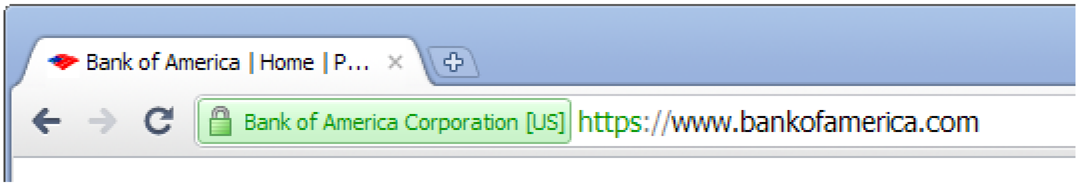

Bank of America | Home | P...
Bank of America Corporation [US]
https://www.bankofamerica.com

Example of the extended validation indicator in Chrome

On most browsers, the extended validation indicator can be found by looking for the name of the organization in the green section of the browser's address bar. You can also click on the indicator to see the website's security information and inspect its digital certificate.

To receive extended validation certification, a website owner has to pass a series of checks confirming their legal identity and authority. In the previous example, extended validation on bankofamerica.com verifies that yes, the website *is* from the actual Bank of America. You can think of this certification as something that ties the domain name of the web address back to some real-world identity.

It'd be wise to share sensitive information with a website only if you trust the organization responsible for the site. So the next time you're about to perform a sensitive transaction, take a moment to keep a look out for the website's security information. You'll be glad you did.

Thing #18:

Evolving to a Faster Web

or, speeding up images, video, and JavaScript on the web

The web today is an amazing visual and interactive stew, teeming with images, photos, videos, and whizzy web apps. Some of the web’s most vivid experiences come from images and videos, from shared photo albums of family vacations to online video coverage from journalists in war zones.

It’s a far cry from the simple text and links that started it all. And it means that every time your browser loads a web page, much more data and complex code needs to be processed.

How much more, and how much more complex? A few astounding statistics:

- Images and photos now make up about **65% of the information** on a typical web page, in terms of bytes per page.
- **24 hours of video** are uploaded to YouTube every minute of the day. (That’s like Hollywood releasing 130,000 new full-length movies every week, though with less popcorn.)
- JavaScript programs have grown from a few lines to **several hundred kilobytes** of source code that must be processed each time a web page or application loads.

So won’t all these gushing floods of data slow down page loads on the browser? Will the Internet clog up and turn to molasses soon?

Probably not. Images and photos became commonplace on the web when computer scientists found ways to compress them into smaller files that could be sent and downloaded more easily. GIF and JPEG were the most popular of those early file-compression systems. Meanwhile, plug-ins were invented to work around the early limitations of HTML so that video could be embedded and played in web pages.

Looking ahead, the <video> tag in HTML5 makes it easy for videos to be embedded and played in web pages. Google is also collaborating with the web community on WebM, an effort to build out a free, open-source video format that adapts to the computing power and bandwidth conditions on the web, so quality video can be delivered to a computer in a farm house in Nebraska or a smartphone in Nairobi.

In the meantime, it's true that web pages with lots of big photos or other images can still be very slow to load. That's why a few engineers at Google have been experimenting with new ways to compress images even further while keeping the same image quality and resolution. The early results? Very promising. They've come up with a new image format called WebP that cuts down the average image file size by 39%.

The engines that run JavaScript code in modern web browsers have also been redesigned to process code faster than ever before. These fast JavaScript engines, such as Google Chrome's V8, are now a core part of any modern web browser. That means the next generation of fabulously useful JavaScript-based web applications won't be hampered by the complexity of more JavaScript code.

Another technique that modern browsers like Chrome use to fetch and load web pages much more quickly is called "DNS pre-resolution". The process of translating a web address into an IP address through a DNS lookup, or vice versa, is often called "resolving." With DNS pre-resolution, Chrome will simultaneously look up all the other links on the web page and pre-resolve those links into IP addresses in the background. So when you do actually click on one of the links on the page, the browser is ready to take you to the new page instantly. Over time, Chrome also learns from past visits so that the next time you go to a web page that you've previously visited, Chrome knows to automatically pre-resolve all the relevant links and elements on the web page.

Someday, browsers might be able to predict, *before* the page loads, not only which links to pre-resolve, but also which website elements (like images or videos) to pre-fetch ahead of time. That will make the web even quicker.

Soon enough, we hope, loading new pages on the browser will be as fast as flipping the pages of a picture book.

Thing #19:

Open Source and Browsers

or, standing on the shoulders of giants

Today's Internet stands on the shoulders of giants: the technologists, thinkers, developers, and organizations who continue to push the boundaries of innovation and share what they've learned.

This spirit of sharing is at the very heart of open-source software. "Open source" means that the inner workings (or "source code") of a software are made available to all, and the software is written in an open, collaborative way. Anyone can look into the source code, see how it works, tweak it or add to it, and reuse it in other products or services.

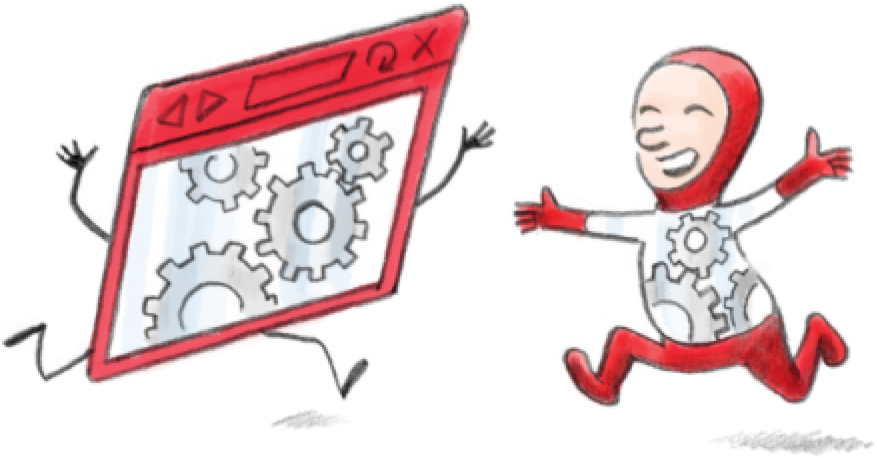

Open-source software plays a big role in many parts of the web, including today's web browsers. The release of the open-source browser Mozilla Firefox paved the way for many exciting new browser innovations. Google Chrome was built with some components from Mozilla Firefox and with the open-source rendering engine WebKit, among others. In the same spirit, the code for Chrome was made open source so that the global web community could use Chrome's innovations in their own products, or even improve on the original Chrome source code.

Web developers and everyday users aren't the only ones to benefit from the faster, simpler, and safer open-source browsers. Companies like Google also benefit from sharing their ideas openly. Better browsers mean a better web experience for everyone, and that makes happier users who browse the web even more. Better browsers also let companies create web apps with the latest cutting-edge features, and that makes users happy, too.

Browsers aren’t the only part of the web that can take the open-source approach. Talk to any group of web developers and you’re likely to hear that they use an open-source Apache HTTP Server to host and serve their websites, or that they developed their code on computers powered by the Linux open-source operating system -- just to name a few examples. The good work of the open source community continues to help make the web even better: a web that can be the broad shoulders for the next generation.

Thing #20:

19 Things Later...

(or, A Day in the Clouds)

...and here we are at Thing 20.

Let's recap.

Today's web is a colorful, visual, practical, nutty, busy, friend-filled, fun and incredibly useful place. Many of us now live a life of cloud computing on the Internet: we read the news, watch movies, chat with friends, and do our daily tasks online with web-based applications right the browser. Web apps let us do that from anywhere in the world, even if we left our laptops at home.

It's all possible thanks to the evolution of web standards like HTML, JavaScript, and CSS, as well as browser plug-ins. And new capabilities in HTML5 are helping developers create the next generation of truly inventive web apps.

What else is taking shape in the clouds?

- It takes a modern browser to make the most of the web's modern features.
- Modern browsers also help protect against malware and phishing.
- Open-source sharing has given us better browsers and a faster, richer, more complex web. And open-source brainpower is making the future of the web even brighter.
- What's in that bright future? 3D in the browser, faster speeds, and sync across all devices,

among other fine things.

- Being an informed citizen of the web requires some self-education -- for instance, learning to control your browser's privacy settings for various types of content including cookies.

- You're also safer on the web when you pay attention to visual cues in the browser, like checking the URLs you're sent to, and looking for an "https://" secure connection or extended validation.

The final takeaways?

Use a modern browser, first and foremost**.** Or try a new one and see if it brings you happier browsing that's better suited to your needs.

The web will keep evolving -- dramatically! Support cutting-edge web technologies like HTML5, CSS3 and WebGL, because they'll help the web community imagine and create a future of great, innovative web apps.

Lastly, try new things. The web is a new and exciting place every day, so try tasks that you didn't think could be done online -- such as researching your ancestry back ten generations, or viewing a real-time webcam image from a climbing basecamp in the Himalayas. You might be surprised by what you find!