

THE  
**TRUTH**  
ABOUT **HTML5**

FOR WEB DESIGNERS • 2012-13 EDITION



SEMANTICS • MICRODATA • SEO • FORMS  
CANVAS • AUDIO/VIDEO • SVG • WEB APPS

BY **LUKE STEVENS**

FOREWORD BY **JOHN ALLSOPP**

**The Truth About HTML5 (For Web Designers)**

Copyright © 2012 Luke Stevens

All rights reserved

*Written and designed by*

Luke Stevens

*Edited by*

Bill Harper

*(Portions of this book have been updated post-edit, so any mistakes are entirely those of the author!)*

*Published by*

Indie Digital Pty Ltd

<http://truthabouthtml5.com>

Spotted an error or typo? Let me know: [luke@itsninja.com](mailto:luke@itsninja.com)

Thanks for reading!

Luke

# TABLE OF CONTENTS

CHAPTER 1	
<b>A Somewhat Sensationlized History Of HTML5</b>	1
CHAPTER 2	
<b>A Basic HTML5 Web Page</b>	21
CHAPTER 3	
<b>Structuring A HTML5 Page</b>	27
CHAPTER 4	
<b>HTML5's Structural Elements</b>	49
CHAPTER 5	
<b>HTML5 Micro-Semantics And Schema.org</b>	75
CHAPTER 6	
<b>HTML5 And SEO</b>	95
CHAPTER 7	
<b>HTML5's Other New Elements</b>	99
CHAPTER 8	
<b>HTML5 Forms</b>	109
CHAPTER 9	
<b>HTML5's Canvas, Gaming And Flash</b>	133
CHAPTER 10	
<b>Audio &amp; Video In HTML5</b>	185
CHAPTER 11	
<b>SVG: The Flash Challenger That Was, Wasn't, &amp; Now...</b>	213
CHAPTER 12	
<b>HTML5 Web Apps, Mobile &amp; What Comes Next</b>	235
CHAPTER 13	
<b>The Future Of Web Design: Performance Based Design</b>	255

## FOREWORD

HTML5 is a mess. It's also one of the most exciting technological advances perhaps ever (a big claim, especially for something I just described as a mess).

There are quite a few books, most of them excellent, on HTML5. Some cover the markup exclusively. Some cover markup and JavaScript APIs. Others still focus on a specific development challenge like games.

This book is a little different. Rather than simply looking at the *what* and *how* of HTML5 (though it does that as well) it endeavors to explain the *why* and *why not* of HTML5.

And it's a passionate, informed, opinionated critique of much of HTML5 to boot.

Along the way you'll learn a great deal about HTML5 markup, and additional HTML5 features such as the new audio and video elements, the Canvas element, the History API, and related features such as SVG.

But hopefully most of all you'll learn to think critically about HTML5 as a tool, and adopt the good parts, for good reasons, and ignore the less than useful parts, for the right reasons as well.

Luke Stevens has written a book all web designers and developers who care about their code should read. So go ahead and read it!

**John Allsopp**

Author, *Developing with Web Standards*

Co-founder of Web Directions

Web evangelist



## INTRODUCTION

Hi. I'm Luke, your average, garden-variety web designer. I've been building web sites for over a decade, use ExpressionEngine as my CMS, and have enjoyed both working in-house and full-time freelancing.

I thought it would be fun to write a short book about HTML5. I thought HTML5 would be simple. I thought writing about it would be straightforward. And I thought the respected voices in the design community would be telling everyone what it is (and what it isn't) simply and clearly, particularly with the plethora of other HTML5 books out there.

I was wrong.

Fortunately this book (and hopefully your experience as a reader!) is infinitely better for it. And I hope once you've read it you'll share my concern about the strange direction basic markup has taken, *and* my excitement for the new HTML5 (and related) technologies that are coming soon to a browser near you. That includes Internet Explorer 10—Microsoft finally, truly gets web standards.

What seemed impossible just a few years ago—a far-fetched, almost utopian ideal of *all* browser vendors, including Microsoft, competing tooth-and-nail to support bleeding-edge web standards—is now a reality. Innovation in web standards is happening at a break-neck speed, and my hope is this book gets you up to speed not only with the fundamentals of HTML5, but with the broader picture of where the web as a whole is heading, especially as we look towards a post-Flash future.

As you make your way through the following chapters, please keep in mind this book is as much of a *critique* as it is an explanation of HTML5. By taking a critical look at *why* things are the way they are, my hope is you save hours by not having to worry about things that don't matter (particularly when it comes to basic markup), and your

eyes are opened to how the HTML5 sausage gets made. It may not always be pretty, but if you spend your days in the trenches building websites, knowing *why* things are the way they are will help guide your design and development decisions in a very direct way.

That said, there's plenty of exciting technology in and around HTML5 too, so be sure not to miss the later chapters on graphics technologies like Canvas and SVG; the state of audio and video in HTML5; and the more developer-oriented HTML5 features that includes a new way of handling something as fundamental as a page request.

(Also note we will be focusing almost entirely on HTML5 as defined by the HTML5 spec, with the addition of SVG, and a few other related initiatives such as Schema.org and WebGL. "HTML5" has become a buzzword which can mean everything from the HTML5 spec itself, to CSS3 and modern JavaScript, to just "cool and new and not Flash". We'll be mostly sticking with the features in the actual HTML5 specification.)

I love the web design community because it's filled with smart, excitable, curious, opinionated folk who will call you on your BS. This is an opinionated book, not a dry explanation of the technology, and I'll be stating my views pretty strongly. I look forward to you doing the same. Passionate, considered debate makes us all smarter. So please, write it up on your blog, send me happy/sad/angry emails ([luke@itsninja.com](mailto:luke@itsninja.com)), talk to me on Twitter ([@lukestevens](https://twitter.com/lukestevens)), or whatever you like.

I look forward to the discussion.

And now I'd like to ask a couple of favors.

First, if you enjoy my writing then please tell your friends, colleagues, Twitter followers, blog readers, and pretty much anyone who will listen about this book. Like a lot of authors, I rely entirely on readers like you to spread the word (and the links). If you can help me out by

spreading the word about this book via good old fashioned word of mouth I'd *really* appreciate it. Thank you.

And second, if you use Google Analytics (and who doesn't?) and want to get more out of it, I'd love you to check out my web app Ninja for Google Analytics at <http://itsninja.com>. Google Analytics is a big, complex beast, but it has the best data on how your web site actually performs, it's just buried deep, deep down. Ninja for GA brings that data to the surface through a simple, elegant interface. It's web analytics for web designers, and I think you (and your clients) will like it. My hope is it will make your own design practice (and your client's sites) more productive and profitable. After all, all the HTML5 in the world won't help you if your conversion rates are lousy and your bounce rates are sky-high. (We'll return to this theme in the final chapter of this book when we look at Performance Based Design.) Check it out: <http://itsninja.com>.



# A SOMEWHAT SENSATIONLIZED HISTORY OF HTML5

## How Architecture Astronauts And The W3C Tried To Kill HTML

Murder is always interesting, so let's start there.

In 1997 the W3C published a Recommendation for HTML 4.0. And two years later it was more or less completed in the form of HTML 4.01. (Don't remember? Well, you were probably too busy worrying about the dreaded Y2K "bug" wiping out civilization.)

And that was pretty much it for plain old HTML.

So what happened between HTML being "finished" in 1999 (in every sense of the word), and HTML5's emergence today?

A long, aborted march to "XMListan".

The W3C published the eXtensible Markup Language (XML) 1.0 spec in 1996 (<http://www.w3.org/TR/1998/REC-xml-19980210>), which they hoped would become a more flexible, machine readable, stricter and more extensible way to mark up documents and other data. And it was soon being used to do just that. But the W3C believed the *web itself* would eventually move to XML.

One of the first baby steps in that direction was XHTML—an XML formulation of HTML 4.

## You Probably Use XML

XML may sound foreign, but if you own or even subscribe to a blog then you're already using it. The RSS or Atom feed blogs generate to syndicate their content is just one form of XML. If you look at the source of an Atom feed, you can see tags such as `<author>`, `<published>`, `<category>` and `<summary>`. These are *specific tags that accurately describe the content they represent*. It's just one example of the “extensible” part of XML that allows machines (parsers, RSS readers and so on) to do interesting things with the content.

Now, imagine a world where we could describe our web pages in a similar way. That was the W3C's plan for the web—that *all* the future content on the web should be described in more accurate terms than just `<div>`s, `<span>`s, `<p>`s and `<h1>`s. And with XML, we could do it.

HTML would still exist as a legacy format. But the future was XML, baby.

## XHTML Is Born, But What Does It Mean?

So if HTML was the past, and XML was the future, how would we get there? With the interim step of XHTML.

By reformulating HTML 4.0 to stick to XML's rules, XHTML was born. And in January 2000, having barely survived the Y2K apocalypse, the XHTML 1.0 spec was adopted as a W3C Recommendation (<http://www.w3.org/TR/xhtml1/>). We were on the road to XMListan.

In early 2002, Jeffrey Zeldman published the landmark XHTML article “Better Living Through XHTML” on A List Apart (<http://www.alistapart.com/articles/betterliving/>), describing XHTML as:

*[T]he standard markup language for web documents and the successor to HTML 4. A mixture of classic (HTML) and cutting-edge (XML), this hybrid language looks and works much like HTML but is based on XML, the web's “super” markup language, and brings web pages*

*many of XML's benefits, as enumerated by the Online Style Guide of the Branch Libraries of The New York Public Library.*

Those benefits enumerated on the The New York Public Library website (<http://legacy.www.nypl.org/styleguide/xhtml/benefits.html>) included:

*The web is moving to XML, a powerfully enabling technology. Writing well-formed, valid XHTML pages is the easiest way to begin this transition. All it takes is learning a few simple rules of XHTML markup.*

Web designers took heed of this call to begin the transition to XML via XHTML. In 2003 Dave Shea wrote a post called “Semantics and Bad Code” ([http://www.mezzoblue.com/archives/2003/08/26/semantics\\_an/](http://www.mezzoblue.com/archives/2003/08/26/semantics_an/)) where he said:

*The move from HTML to XML requires a huge shift in developer mindset. There are a lot of obstacles to overcome yet, not the least of which being solid browser support. We've only started down the road, and XHTML is how we'll get there.*

Shea's view was a popular one at the time, and certainly reasonable given our faith in the experts in the W3C.

But we never made it to XMListan. The car ran out of gas, the wheels fell off, and the engine exploded about two blocks down the road.

## **Draconian Error Handling, Or Why Don't I Just Punch You In The Face?**

Those of you building web sites back in the early '00s may remember how important it was to have a *valid* web page. People even put dinky little “Valid XHTML” badges on their sites to show off just how forward-thinking they were. (They now put equally silly HTML5 badges on blogs—and books.) Design nerds would even run *other people's* markup through the HTML validator, and write a snarky blog post or email if it failed. (Back then there was no Twitter to bitch publicly in 140 characters.)

Yes, having valid HTML is a good thing. But as web designers adopted XHTML it became—in theory, if not practice—life or death. If you had so much as single error in your XHTML, your browser would reach out and *punch you in the face*.

## Okay, Not Really. But We COULD Punch You In The Face

Well, it would if you set up your server to tell the browser to adopt XML's strict XHTML parsing rules (as Mark Pilgrim described in 2003: <http://www.xml.com/pub/a/2003/03/19/dive-into-xml.html>), which hardly anyone did. Internet Explorer, right up to and including version 8, didn't even support these strict XHTML parsing rules. (Ironically, IE9 now does, just as everyone stopped caring.)

Why didn't anyone do it? Because they didn't want to inflict the “draconian error handling” on their users (or themselves). And it really was draconian—one invalid character, such as “&” instead of “&amp;”, would generate a fatal error that destroyed the entire page. And as a user, all you got was a hideous error message—no content, no nothing.

In light of this, the web standards community adopted the *theory* of XHTML without its harsh realities (or true XML nature), preferring to stick with the warm, cuddly and vastly forgiving HTML parsing from the early days.

XHTML turned out to be a baby step towards a baby step. What should have been the first move towards a strict XML formulation of the web, where we could use more descriptive (i.e. semantic) tags, was just a step towards stricter, old-style HTML. It was two steps forwards, one step back—back to the HTML the W3C had declared finished, and was hoping to make obsolete.

## XHTML Still Meant Better HTML

Nevertheless, XHTML gave the web standards community something to, well, standardize on. It allowed everyone to be a bit more serious, and dare I say *professional*, about the markup we were writing. As Jeffrey Zeldman wrote on his blog in 2009 (<http://www.zeldman.com/2009/07/07/in-defense-of-web-developers/>):

*XHTML's introduction in 2000, and its emphasis on rules of construction, gave web standards evangelists like me a platform on which to hook a program of semantic markup replacing the bloated and unsustainable tag soup of the day. The web is better for this and always will be, and there is much still to do, as many people who create websites still have not heard the call.*

For much of the '00s, websites built with web standards continued using XHTML. Designers got serious about separating presentation from content, and tried to write more semantic markup. Using XHTML also triggered standards mode on the major browsers of the time. All good things.

But in the W3C's grander scheme of things, XHTML ultimately proved to be a bit of a stepping stone to nowhere.

## **But The Crazy Had Only Just Begun**

XHTML served a useful purpose for web standards—albeit not the one originally intended. But now we step into the mad, mad, mad world of XHTML **2.0**.

While we were all happily using and advocating XHTML in web standards land (though some stuck to HTML 4.0), the W3C was working on XHTML 2.0. Sounds like a harmless update of the 1.0 spec, right?

It wasn't.

XHTML 2.0 was day zero for the web. It wasn't backward compatible with HTML, or even XHTML 1.0. It was a *whole new thang*.

And nothing was safe.

Among the list of sweeping changes, plain old forms would be replaced with fancy XML-style XForms. Even the `<img>` element was on the chopping block at one point, as the W3C re-envisioned the web as a more XML-ified place.

In an April 2011 blog post on software development, Joel Spolsky described what he calls “Architecture Astronauts” (<http://www.joelonsoftware.com/articles/fog0000000018.html>):

*When you go too far up, abstraction-wise, you run out of oxygen. Sometimes smart thinkers just don't know when to stop, and they create these absurd, all-encompassing, high-level pictures of the universe that are all good and fine, but don't actually mean anything at all.*

*These are the people I call Architecture Astronauts.*

And XHTML 2.0 was a classic case of Architecture Astronauts at work.

Here’s how Bruce Lawson, HTML5 evangelist for Opera and author of “Introducing HTML5” (New Riders, 2010) describes it ([http://news.cnet.com/8301-17939\\_109-10281477-2.html](http://news.cnet.com/8301-17939_109-10281477-2.html)):

*XHTML 2 was a beautiful specification of philosophical purity that had absolutely no resemblance to the real world.*

As far as HTML was concerned, this is what the W3C—the custodians of the language that underpins much of our relationships, business, and government in the 21st century—worked on from 2002-2006 over 8 drafts. Not only would it have broken backwards compatibility, it would also have sent all the talk of “forward compatibility” and “future-proofing” in the web standards community up in smoke. (You can read more about XHTML 2.0 in Wikipedia: [http://en.wikipedia.org/wiki/XHTML#XHTML\\_2.0](http://en.wikipedia.org/wiki/XHTML#XHTML_2.0).)

## **XHTML 2.0: Unloved And Alone**

While the W3C toiled away on XHTML 2.0, what did web authors, standards advocates, and browser vendors think of it?

Not much.

There was zero interest in implementing it. Even members of the working group were deeply unhappy with it. (See Jeffrey Zeldman's thoughts on XHTML 2.0 in 2003 under "XHTML 2 and all that": <http://www.zeldman.com/daily/0103b.shtml>.)

What was dopey about XHTML 2.0 wasn't so much the spec itself (which would be fine if we could go back in time and rebuild the web from scratch). It was the idea you could do something as revolutionary as breaking backwards compatibility with millions of existing documents and create a whole new tier for the web. But that was the path the W3C set themselves on way back in 1998 (see it for yourself in "Shaping the Future of HTML" <http://www.w3.org/MarkUp/future/>).

But what if the next evolution of HTML was just that—evolutionary, rather than revolutionary? One that built on the world as it was, and not some utopian world we could only hope for?

## HTML5: A New Hope... We Hope

HTML5 began as a reaction to the W3C's descent into markup madness. The problems with the W3C's direction had not gone unnoticed.

In 2004, the so-called "Web 2.0" movement took off in a big way, and web *applications* became a big deal. The web was no longer just a collection of text and images on pages connected through links. It was becoming a platform for applications that could run anywhere, OS be damned.

Compared to the '80s and '90s, when your OS determined what applications you could use, running applications through a browser on any OS was a revolutionary idea.

No one really predicted this (certainly not the W3C), which isn't surprising when you think how bad we are at predicting the future in general. (Where *is* my flying car?) We're much better at reacting and evolving when the future arrives, which is what some people suggested we do with HTML.

In 2004, members representing Opera and Mozilla (with Apple "cheering [from] the sidelines", as Ian Hickson recalls: <http://www.webstandards.org/2009/05/13/>

[interview-with-ian-hickson-editor-of-the-html-5-specification/](#)) presented an alternative to the W3C—a spec focused on web applications. (See the original “Position Paper” here: <http://www.w3.org/2004/04/webapps-cdf-ws/papers/opera.html>.)

## The W3C Says Go To Hell

HTML needed to adapt to the future of web *applications*, rather than a utopian world of perfectly marked-up XML-ified web *pages*. So this new group suggested an alternative direction for HTML based on backwards compatibility. No more draconian error handling (the one-error-and-you’re-dead problem of XHTML as XML). New features for web applications. And an open process, which was in stark contrast to the way the W3C operates.

Essentially, their philosophy was that HTML was here to stay, and so we should concentrate on evolving it. (This may sound completely obvious now, but back then it wasn’t a view shared by the W3C.)

Anyway, the group pitched their ideas to the W3C, and the W3C told them to go to hell. (Actually, they only lost by two votes—11-8 against. But this *is* the somewhat *sensationalized* history of HTML5.)

With the W3C being less than accommodating, those interested in evolving HTML and adding features for web applications, and who were backed by (and worked for) the browser vendors, decided to press on and work outside the W3C. They formed the Web Hypertext Applications Technology Working Group (WHATWG), and set up shop at [whatwg.org](http://whatwg.org) in June 2004.

## The WHATWG Is Born

And so the WHATWG was born. Here’s how Hickson explains it all (<http://www.thechromesource.com/interview-html5-standards-author-ian-hickson/>):

*So [after the W3C rejection] we opened a mailing list called the WHATWG to continue work on Web Forms 2.0 in public, and later that year started a new draft called Web Applications 1.0 into which*

*we put many features aimed at writing Web apps, including a new version of HTML that we jokingly called HTML5, and a bunch of other features that later became Web Storage, Web Sockets, Server-Sent Events, and a variety of other specs. [...]*

*Later, around 2006 or 2007, the W3C basically realized they had made a mistake, and they asked if they could work on HTML5 as well, so we renamed Web Applications 1.0 to HTML5, and the WHATWG and the W3C started working together. Web Forms 2.0 got merged into HTML5, and most of the bits of HTML5 that weren't really HTML got split out into separate specs.*

It's ironic, isn't it? The establishment (the W3C) was the utopian revolutionary, and the rebel outsiders (the WHATWG) were fighting for incremental conservatism. Go figure.

## It's A Whole New World

It's worth noting several points here:

- The W3C failed dramatically at maintaining HTML (which is kind of scary when you think about it).
- Web standards are incredibly haphazard. There was—and is—no unifying vision of “HTML5”. It was just a bunch of separate specifications bundled up and given the name “HTML5”, and those specifications only came about as a reaction to the W3C's failures.
- Big, bold ideas like the march to XML for the web—which had many people excited a decade ago—can fade to nothing. We should learn from this, and retain some skepticism towards big, bold ideas—including some of the changes in HTML5.
- The balance of power now rests with the browser vendors.

In truth, the balance of power has *always* rested with the browser vendors. If they don't implement something, *by definition* it's a non-starter. As Hickson says (<http://www.webstandards.org/2009/05/13/interview-with-ian-hickson-editor-of-the-html-5-specification/>):

*The reality is that the browser vendors have the ultimate veto on everything in the spec, since if they don't implement it, the spec is nothing but a work of fiction. So they have a lot of influence—I don't want to be writing fiction, I want to be writing a spec that documents the actual behavior of browsers.*

*Whether that's too much, I don't know. Does gravity have too much influence on objects on earth? It's just the way it is.*

Nevertheless, the fact an independent standards body—*our* independent standards body—failed miserably is more than a little concerning.

## To HTML5 And Beyond!

To cut a long story short, the WHATWG kept working on their own vision of evolving HTML—the *only* vision of evolving HTML. And in 2006 Tim Berners-Lee, father of the World Wide Web and Director of the W3C (read more about him here: [http://en.wikipedia.org/wiki/Tim\\_Berners-Lee](http://en.wikipedia.org/wiki/Tim_Berners-Lee)), sucked it up and announced the W3C would work with the WHATWG, saying (<http://dig.csail.mit.edu/breadcrumbs/node/166>):

*The attempt to get the world to switch to XML, including quotes around attribute values and slashes in empty tags and namespaces all at once didn't work.*

Berners-Lee left the door open to switching to XML by saying “all at once”. But in reality it looks very much like “The attempt to get the world to switch to XML... didn't work.”

And that's fine. We need big ideas and bold directions to try and work towards, and if they don't work out, so be it. Sometimes good ideas just don't happen.

With the WHATWG having so much momentum (and the backing of the browser vendors), the W3C had no choice *but* to work with them on HTML5. In 2007 the W3C formed a group that worked work with the WHATWG on developing HTML5. And in January 2008 the W3C released their first HTML5

Working Draft (<http://www.w3.org/TR/2008/WD-html5-20080122/>), adopting the work the WHATWG had been doing for several years.

## HTML5 Is The New Black Or Hotness Or Something

By the late '00s web technologies were exciting again, and after years of stagnation and dead ends we finally reached a point where the bowels of innovation were loosened. (That's a horrible image—sorry.)

Now in the early '10s, things are looking even better. In fact, there's a veritable Cambrian explosion of web technology taking place. Google, Mozilla, Apple and Microsoft are competing to make the best *standards compliant* browser (with new versions coming thick and fast). There's a whole bunch of new and interesting technology around. And web developers, designers, software companies and app developers are all interested in the new and shiny tech in and around HTML5.

To think browser makers—including Microsoft—are now trying to outcompete and even out-market each other with their *web standards support* is pretty incredible. It wasn't that long ago (late '90s) that we faced the threat of them all going their own non-standard ways. Hats off to all involved.

## Is HTML5 Hype, Substance, Or Both?

But back to the HTML5 specification. Two questions:

1. What exactly *is* HTML5?
2. Who's in charge, now there's a (decidedly uneasy) working relationship between The Establishment (the W3C) and The Rebels (the WHATWG)?

Let's deal with what HTML5 *is* first. There's:

- HTML5, the all-encompassing marketing buzzword
- HTML5, the bit that's actually about HyperText Markup
- HTML5, the new functionality available through JavaScript for web applications

- HTML5, the behind the scenes stuff that's really important and documents a whole lot of stuff browsers actually do (but you're probably not interested in).

All this from a technical specification that runs for hundreds of pages.

For us web designers, HTML5 is currently a confusing mix of hype *and* substance, which we'll try to sort through in the coming chapters.

In many ways HTML5 is, to put it bluntly, a mess. But it's the most ordered mess we've had in a long time. (For instance, a big part of HTML5 is written for browser vendors to ensure implementations are consistent and we can trust all browsers to do the same thing. And that's never been done before.)

Perhaps the biggest problem is everyone thinking that if HTML5 is cool, then *all* of it (at least according to the web design community) *must* be great, and we should adopt it post-haste without too much critical thought. And that's something I'm keen to dispel in the rest of the book.

## Hixie Or Bust

As I write this, both the WHATWG and W3C versions of the HTML5 spec (the differences between the two are minor) are edited by one person: Ian Hickson.

HTML is now essentially in the hands of one man.

The W3C's working groups tried building consensus, and got absolutely nowhere with HTML. It was closed, but democratic. The WHATWG, on the other hand, has an open process, but with an editor-has-the-final-say approach.

And that editor is Ian "Hixie" Hickson.

Hickson helped start the WHATWG when working for Opera, and now works full-time for Google developing the HTML5 spec. Currently, he is the HTML5 (and now just "HTML") editor for *life*. Theoretically, the browser makers can veto him or kick him out at any time, but that seems highly unlikely. This has not gone unnoticed in the community, and is (rightly, in my opinion) a cause of some concern.

It's a classic "glass half-full/glass half-empty" situation. If Hickson flat out refuses an idea (which is known to happen), then having a single person in charge may seem like utter madness. But for those who saw the W3C's democratic processes get nowhere with XHTML 2.0, having someone who can take the reins, push things along, and actually make decisions would seem wonderful.

Of course, this invariably polarizes people.

Here's John Gruber of Daring Fireball fame (<http://daringfireball.net/linked/2009/07/01/hickson-codecs>):

*Let it be said that Ian Hickson is the Solomon of web standards; his summary of the situation is mind-bogglingly even-handed and fair-minded.*

And here's Kyle Weems, creator of the CSSquirrel comic, who has been following HTML5's development for several years (<http://twitter.com/#!/cssquirrel/status/58559284224589824>):

*Also... why oh why is @hixie still the editor for any world-altering spec like HTML anymore? Ego doesn't even begin to describe his antics*

As you can see, Hickson has his fans and his detractors.

I imagine editing a spec the size of HTML5 for as long as he has, with all the controversy that surrounds it, would be a pretty thankless task. But Hickson seems to go about it in a cheerful, dispassionate way.

If there's one overarching theme here, it's this: pragmatism rules.

The W3C had the "pure" spec of XHTML 2.0, and failed—it wasn't pragmatic. It also had its rules, membership, and democratic processes, but was mired in politics and failed (with HTML at least)—it wasn't pragmatic.

The WHATWG put an editor in charge, and while this approach terrified and/or infuriated some people (including me from time to time, as you'll soon see), it

was pragmatic (as was their approach to the spec). It got things moving (and, more importantly, shipping). And as long as it remains pragmatic it's probably how the WHATWG will stay.

## XHTML 2.0 Is Dead And Everyone Is Happy

So what happened to XHTML 2.0? It was pronounced dead after being taken off life support in 2009 (<http://www.w3.org/2009/06/xhtml-faq.html>). I hear the death of XHTML 2.0 will soon be fictionalized in an upcoming episode of "Law & Order: Web Standards Unit".

And what about XHTML 1.0 and its various flavors? Considering it's essentially just HTML, it will keep working pretty much forever. (There's actually a continuing XML serialization of HTML5 called XHTML5, but the chance of you actually needing to use it is practically zero.)

## HTML5, err HTML, wait... HTML.next?

To show how things have come full circle with the HTML spec, the WHATWG declared in January 2011 that their HTML5 spec would be a "living standard" and renamed it to just "HTML". (See the announcement here:

<http://blog.whatwg.org/html-is-the-new-html5> and their rationale here:

[http://wiki.whatwg.org/wiki/FAQ#What\\_does\\_.22Living\\_Standard.22\\_mean.3F.](http://wiki.whatwg.org/wiki/FAQ#What_does_.22Living_Standard.22_mean.3F.))

And what of the future of HTML? The WHATWG insist they—and particularly Hickson—will maintain the HTML spec as a "living standard" indefinitely, while the W3C are sticking with the snapshot process, and have started accepting ideas for what they're unofficially calling "HTML.next" (see some of the ideas here: <http://www.w3.org/wiki/HTML/next>).

(A W3C member gave a personal presentation that captures the differing approaches to the future of HTML quite nicely: <http://www.w3.org/2010/11/TPAC/HTMLnext-perspectives.pdf>.)

Will the W3C come up with another pie-in-the-sky path to nowhere (echoing 1998's "Shaping the Future of HTML" workshop <http://www.w3.org/MarkUp/>

future)? Will they try to work with the WHATWG, or fork HTML5 and do their own thing? Who knows. Some have been asking if the W3C should even exist.

## Should We Just Kill Off The W3C Altogether, Or Embrace It?

In September 2011, a debate broke out about the purpose of the W3C, and three broad views emerged: reform, destroy, and embrace.

Before we get to those three views, let's consider why debate about the W3C is still continuing, just as it seems to have its house in order, having adopted the WHATWG's successful HTML5 specification.

In short, it's because the world kept turning. The WHATWG began their work on what became HTML5 in the mid-2000s, and the details of HTML5 (and related specifications) are still being nussed out in the 2010s. Mobile is exploding, "apps" are taking us back to the platform-specific software world of the 90s, and standards development is still slow, even in this new wow-stuff-is-actually-happening environment we now enjoy.

Can the web keep up in the face of resurgent, platform-specific app development? Has the W3C outlived its usefulness, or is it now finally back on track after years in the wilderness? Here are three perspectives, all from September 2011:

### Reform

In "Things the W3C Should Stop Doing" (<http://infrequently.org/2011/09/things-the-w3c-should-stop-doing/>), Alex Russell, who works for Google on Chrome, argues the W3C needs to drop all its XML and enterprise stuff, and refocus solely on the web. Essentially, drastic reform can save the W3C from irrelevance.

*The time has come for the W3C to grab the mantle of the web, shake off its self-doubt, and move to a place where doing good isn't measured by numbers of specs and activities, but by impact for web developers.*

## Destroy

In “Web Technologies Need an Owner” (<http://joehe Witt.com/2011/09/22/web-technologies-need-an-owner>), Joe Hewitt, who worked on early versions of Firefox, created Firebug, and was responsible for the iPhone Facebook app, argues the web is just another platform, but without anyone taking responsibility for it (unlike Windows, Android and iOS).

*Let's face facts: the Web will never be the dominant platform. There will forever be other important platforms competing for users' time. To thrive, HTML and company need what those other platforms have: a single source repository and a good owner to drive it. A standards body is not suited to perform this role. Browser vendors are innovating in some areas, but they are stalled by the standards process in so many areas that is impossible to create a platform with a coherent, unified vision the way Apple has with Cocoa or the way Python has with Guido.*

Therefore we should, as Hewitt tweeted (<https://twitter.com/joehe Witt/status/116292923288592384>):

*[D]issolve the W3C, and run the web like an open source project. No more specs, just commits. Does Linux need a standards body?*

## Embrace

Finally, in “The web is a different problem” (<http://www.webdirections.org/blog/the-web-is-a-different-problem/>) John Allsopp, long standing web evangelist, writer, and speaker, argues that while standards development certainly stalled in the 00s, we’ve seen an “explosion of innovation at the browser level” in the last few years, particularly with CSS3 and more modular specs, and are we really now going to throw the baby out with the bathwater?

*So, to put it bluntly, I think the problem is overstated. We seem to have arrived at an approach that both enables the exploration and implementation of novel features in browsers, which are also widely adopted across browsers. [...]*

*[But] the web is a different problem. It makes little if any sense to compare innovation of the web ecosystem with that of iOS, Android or other platforms. The web faces challenges far far greater (and has goals far more important). [...]*

*So, rather than generally criticising the W3C, or going so far as calling for its dissolution, we should focus on how well in many ways it has done an almost impossible task—getting companies which are fierce commercial rivals to sit down, work together and agree on core technologies they will each, and all, implement, even while at the same time, these same competitors are involved in significant legal conflicts with one another.*

Whatever we may wish for, sheer inertia is likely to see the W3C maintain its role as the home of web standards development in the coming years (for better or worse), especially now it has brought the WHATWG and HTML5 inside the W3C tent.

## How Does New Stuff Get Added To HTML5 Now?

How will HTML5 evolve from here on out? How will the WHATWG implement new HTML features in their “living standard”? They say new HTML features should first appear in browsers (experimentally at least), and *then* be codified into the spec, assuming there’s a reasonable use case for them and the editor approves. (See the WHATWG FAQ for more: [http://wiki.whatwg.org/wiki/FAQ#Is\\_there\\_a\\_process\\_for\\_adding\\_new\\_features\\_to\\_a\\_specification.3F.](http://wiki.whatwg.org/wiki/FAQ#Is_there_a_process_for_adding_new_features_to_a_specification.3F.))

This means the HTML spec will capture features as they emerge, rather than dictate new features from scratch—a somewhat odd stance given the amount of innovation the WHATWG did in the HTML5 spec before *any* browser implementation.

How long will the WHATWG/W3C relationship last? Your guess is as good as mine. Hickson has been openly hostile to the W3C’s process at times (<http://lists.w3.org/Archives/Public/www-archive/2012Jan/0032.html>), and his

decisions and refusals continue to be a source of considerable friction on the W3C mailing lists.

At the end of the day, either party can dream up all the specs they like. What really matters is what the browser vendors choose to implement. As far as HTML is concerned, the WHATWG's extremely close relationship with the browser vendors means they'll probably be calling the shots for the foreseeable future.

So, after all that, we're back to HTML. And that wraps up our somewhat sensationalized (and highly condensed) history of HTML5. Or HTML. Or... you get the idea.

## TL;DR

In summary, the W3C tried to kill HTML and took us on a decade-long journey to nowhere; some people from browser vendors formed a group interested in web apps and evolving HTML's forms; they worked outside the W3C on what became HTML5; the W3C realized they were screwed and agreed to use their work; browser vendors are implementing it (or their existing implementations of certain features have been standardized); web standards have become a *Microsoft* marketing buzzword; hell has not frozen over.

## What We'll Be Focusing On

HTML5 is a massive specification, filled with mind-numbing detail for browser vendors.

But that detail is actually the best thing about it. Removing the implementation ambiguities has led to more predictable behavior, which is good news for designers and developers alike. (Before that, browser vendors were looking over each other's shoulders to see how parts of the spec were interpreted.)

It's not sexy work, but rather years of careful documentation and clarification by the WHATWG that we can all be grateful for.

The other parts of HTML5 very much reflect its origins as Web Applications 1.0 and Web Forms 2.0. We'll touch on the web app stuff in chapter twelve, and look at the web forms in chapter eight.

As designers, the biggest point of interest are the changes and additions to the actual markup side of HTML. And that's what we'll focus on: semantics, forms, graphics, and audio/video.

We'll also touch on the new features for web apps in HTML5, which we'll hopefully see in our Content Management Systems sooner rather than later.

Most importantly though, we'll be looking at the *ideas* in markup and the practical—sometimes critical—dos and don'ts of HTML5 along the way.

Let's jump in, and look at how we start a document in HTML5.



# THE TRUTH ABOUT A BASIC HTML5 WEB PAGE

## A Doctype For Every Occasion (And The Other Bits)

Let's start with the first line of a web page. It's now just:

```
<!doctype html>
```

That's it. Short, memorable, and triggers standards mode in all major browsers (including IE6). It's also case insensitive.

In HTML5 the opening `<html>` tag has also been simplified to:

```
<html lang="en">
```

Browsers will cope without the `lang` attribute, but it's good practice to specify the page's primary language—especially for non-English pages. (See this helpful article on declaring languages in HTML5: <http://nimbupani.com/declaring-languages-in-html-5.html>.)

Next comes the `<head>` tag, which will contain our `<title>`, `<meta>`, CSS and JavaScript tags as per usual. You don't actually need to specify `<head>` tags if you want to be ultra minimal (see Bruce Lawson's minimal HTML5 document discussion here: <http://www.brucelawson.co.uk/2010/a-minimal-html5-document/>), but we will.

Inside our `<head>` tags we have:

```
<meta charset="utf-8">
```

This specifies the character encoding for the page. Again, it's been reduced to the simplest form possible in HTML5. You should *always* specify this for security reasons (there's a technical discussion here: <http://code.google.com/p/doctype-mirror/wiki/ArticleUtf7>), and it should come before the `<title>` tag.

```
<meta name="description" content="My HTML5 Website">
```

This hasn't changed. Google and other search engines sometimes use this tag in their search results pages, but not for rankings. (You can forget all about `<meta content="keywords">` though. Search engines have been ignoring it for years. We'll look at markup and SEO in chapter six.)

For more on meta tags Google does understand, see: <http://www.google.com/support/webmasters/bin/answer.py?answer=79812>.

The `<title>` tag hasn't changed.

To link CSS and JavaScript files, we can just use:

```
<link rel="stylesheet" href="styles.css">
```

And:

```
<script src="myscript.js"></script>
```

There's no need to specify `type="text/css"` or `type="text/javascript"` anymore—the browsers assume it anyway.

We can start using these techniques now. There's no harm in them, they just make it simple enough to start writing our documents from memory. (The old techniques will continue to work though—probably forever.)

So, a basic HTML5 page (with basic body content) looks like this:

```

<!doctype html>
<html lang="en">
<head>
  <meta charset="utf-8">
  <meta name="description" content="My HTML5 Website">
  <title>My HTML5 page</title>
  <link rel="stylesheet" href="mystyles.css">
  <script src="myscript.js"></script>
</head>

<body>
  <h1>My HTML5 Page</h1>
</body>
</html>

```

As you can see it's pretty much what we're used to, just simpler.

## Formatting Changes In HTML5

A few things to note about how we write HTML in HTML5:

- **Quotes are optional.** You no longer need to quote attribute values, so you can write `<meta charset=utf-8>` or `<div class=myclass>` if you like. Personally I prefer quoting values, but HTML5 leaves it up to you.
- **It's case-insensitive.** You can write your markup in upper or lowercase, or even a mix like `<DiV ClAsS=VaLuE>` if you really hate your coworkers and/or feel nostalgic for YoUr WaCkY MySpAcE days.
- **Closing slashes are optional.** You no longer need to close standalone tags with a closing slash (e.g. `<meta charset=utf-8 />`). As you probably guessed, this was a relic of the move to XML. Likewise, `<br>` and `<br />` are both perfectly valid—it's up to you.

If you're a stickler for XHTML's stricter syntax (always writing in lowercase, quoting attribute values and closing standalone tags), you can keep doing it—it will always be happily supported.

## What About A HTML5 Shim And CSS For The New Elements?

HTML5 introduces new elements such as `<nav>`, `<header>`, `<article>`, `<section>`, and so on. These sound fine in theory, but are terrible in practice.

To support these elements in IE6-8, others suggest you include a small script that tells IE6-8 these elements exist and to use whatever styles you specify for them (it will leave them unstyled otherwise). I don't recommend using these new elements, so we don't need the HTML5 shim. (If you really want to use them, here's the code to do it: <http://code.google.com/p/html5shiv/>. But seriously, don't use the new elements. You'll thank me later.)

You also need to set the new elements to `display: block;`, as shown in this HTML5doctor.com boilerplate: <http://html5doctor.com/html-5-boilerplates/>. Again, don't use these elements. (You'll see why in the next two chapters.)

## What About The HTML5 Boilerplate And Modernizr?

If you want an everything-and-the-kitchen-sink boilerplate for new HTML5 pages, check out <http://html5boilerplate.com/> and the markup documentation <https://github.com/paulirish/html5-boilerplate/wiki/The-markup>. (There's more documentation in the wiki.)

While I appreciate the effort they've put into the HTML5 Boilerplate, if you're just finding your way with HTML5 it's pretty intense. I prefer to start simple and work with my own bare-minimum approach. But if you prefer the start-with-everything-and-delete-what-you-don't-want approach, the HTML5 Boilerplate may be right up your alley.

Modernizr (<http://www.modernizr.com/>) is a handy script for detecting support for HTML5 and CSS3 features. (It doesn't *add* support, it only *detects* it.) It's become a staple for designers who live on the bleeding edge and experiment with new features, so if that's what you're interested in check it out. (We'll talk more about Modernizer, and the merits of *feature* detection rather than *browser*

detection, when we look at HTML5's web application features in chapter twelve.)

Well, that was easy. Almost *too* easy. Now let's take a big left turn into the proverbial ditch that is the new structural tags.



# THE TRUTH ABOUT STRUCTURING A HTML5 PAGE

## New Structural Elements - This Isn't Going To End Well (Plus, Controversy!)

One of the most common tasks web designers do is mark up page structure, which usually consists of a header, footer, navigation, sidebar and content area. It's the sort of thing you can probably do blindfolded and handcuffed to your chair after being spun around for five minutes.

HTML5 introduces a handful of new elements to help us define the structure of a given web page, such as `<section>`, `<article>`, `<nav>`, `<aside>`, `<header>`, and `<footer>`.

We shouldn't use them. They were made up on a whim by (probably) *one guy* in 2004 and even *he* seems to have forgotten what their purpose is.

If that's all you needed to know, great. Keep using `<div>`s with meaningful class and ID names, and appropriate `<h1>`-`<h6>` headings. They'll be valid forever (more or less), and you're not missing out on anything.

However, I suggest using some non-HTML5 features when marking up documents, such as ARIA attributes for blind and sight-impaired users and microdata schemas (when appropriate) for search engine results. (We'll talk more about these in later chapters.)

Nevertheless, we'll tackle these new elements in depth because everyone gets them wrong. And we'll set the record straight on how they found their way into the spec and their real intended purpose, which involves a radically different way of structuring your pages.

## A Little Taste Of Pain

Here are just some of the problems these new structural elements introduce:

- They give terms web designers already use (such as header and footer) new uses, while claiming to be just doing what web designers are already doing.
- They introduce a new method of structuring documents that's vague, complicated, and unnecessary.
- They seriously hurt accessibility for some users (specifically those using IE6, IE7, and even IE8 with JavaScript switched off).
- They introduce broad, unclear, poorly-defined use cases that will make web standards harder to learn (and harder to teach).

These are serious problems that hurt, rather than help, web standards. Markup should be lightweight, easy to learn, and easy to apply. It should *not* require mental gymnastics to try and work out what to use where.

But these new structural tags have created a strange, quasi-religious experience where you have to consult the high priests (the HTML5 gurus) for their interpretation of vague religious texts (the HTML5 spec) just to mark up a darn web page.

“But, but... these elements are in the official HTML5 spec! Surely there *must* be a good reason for them?”

Read on...

## Where Did These Elements Come From?

Quiz question: How were these elements added to the HTML5 spec?

- a. Experts considered various use cases, weighed up various options and alternatives, and after extensive consultation and careful deliberation included the most important ones.
- b. The community of web designers and HTML authors (such as you and me) cried out for certain elements to enable particular functionality, and after much discussion the community came up with a shortlist of necessary elements.
- c. A scientific, research-based approach was taken, where markup patterns were studied “in the wild” and codified into a bunch of new elements.
- d. Some markup wonks thought they’d be a good idea and threw them in the spec 7+ years ago.

And the answer is... (d).

“But I read in [insert HTML5 book of your choice here] that it was more like answer (c). The WHATWG studied real-world usage of ID and class names, and that’s how they came about!”

We’ll get to that.

I was intrigued about who added these elements, when they added them, and why. So I put those questions to HTML5 spec editor Ian Hickson, and here’s his reply (reproduced with permission):

*Me and other WHATWG contributors [added them], [in] 2004ish, because they were obvious elements to add after seeing how authors used HTML4. We later (late 2005 early 2006) did some objective research to find out what the top ten HTML classes were and it turned out that they basically exactly matched the elements we had added, which was convenient.*

You may have read about this “objective research” in other HTML5 books, in talks on HTML5, or in blog posts about these new elements. But almost everyone fudges the history. Sometimes they say the research came first—it didn’t. Sometimes it’s just implied the research came first, which is still a sin of omission.

(Actually, according to the research in question—<http://code.google.com/webstats/2005-12/classes.html>—the major finding was that around 90% of the billion pages

sampled had no classes at all. If Hickson and the WHATWG truly followed the research here, they would have abolished classes altogether!

So if these elements didn't come about from research, where did they come from?

Exploring the dark recesses of the (thankfully public) WHATWG mailing list, I found Hickson first mentioning these elements in November 2004, when he discussed block level elements listed on his whiteboard. (See:

<http://lists.whatwg.org/htdig.cgi/whatwg-whatwg.org/2004-November/002329.html>.)

In the same week he said “[W]hat I'm thinking of doing is [adding] section elements [that] would be: `<body> <section> <article> <navigation> <sidebar>`”. (You can see the full email here: <http://lists.whatwg.org/htdig.cgi/whatwg-whatwg.org/2004-November/002362.html>.)

Of course, somewhere along the way `<navigation>` became `<nav>`, and `<sidebar>` became `<aside>`.

So these new, major structural elements that everyone is trying to get their heads around were probably included because Hickson jotted them down on his whiteboard in 2004. They actually serve a much broader purpose for “sectioning” (which we'll get to shortly). But it's worth establishing how they wound up in the spec, and how arbitrary they are.

In chapter one we saw that XHTML 2.0 failed for being absurdly ambitious. In HTML5 we instead get a few semantic elements the editor drew on a whiteboard years ago on a whim, with some input from a handful of fellow WHATWG members of the time.

## Who Cares?

“Well, who cares?” you may think. “If the research ultimately supported using these elements, then what's the big deal?”

The problem is Hickson was, in my view, being a bit cheeky when he said these new elements “exactly match the elements we ... added.” While they share the

same *name* as elements commonly used, the spec describes their use in *very* different ways to what the web designers and authors would be familiar with. And for a standard these web designers and authors are supposed to use, that's a big problem.

What happens when you take terms people use, redefine how they should be used (and even give them multiple uses), and then tell those same people not to worry because the terms are exactly what they're already using? You put them on a one-way trip to confusion city.

## The Contradiction At The Heart Of HTML5's New Elements

HTML5 is supposedly about codifying what we're already doing, or "paving the cowpaths". When it comes to these new tags and marking up a basic template, they suggest you can just replace your current `<div>` structural tags with the new tags (e.g. replace `<div id="header">` with `<header>`), and you're done.

That was certainly the implication in the December 2007 ALA article "A Preview of HTML 5" (<http://www.alistapart.com/articles/previewofhtml5>), and the idea has been repeated in books and blog posts since, usually with a graphic like the underlying one here:

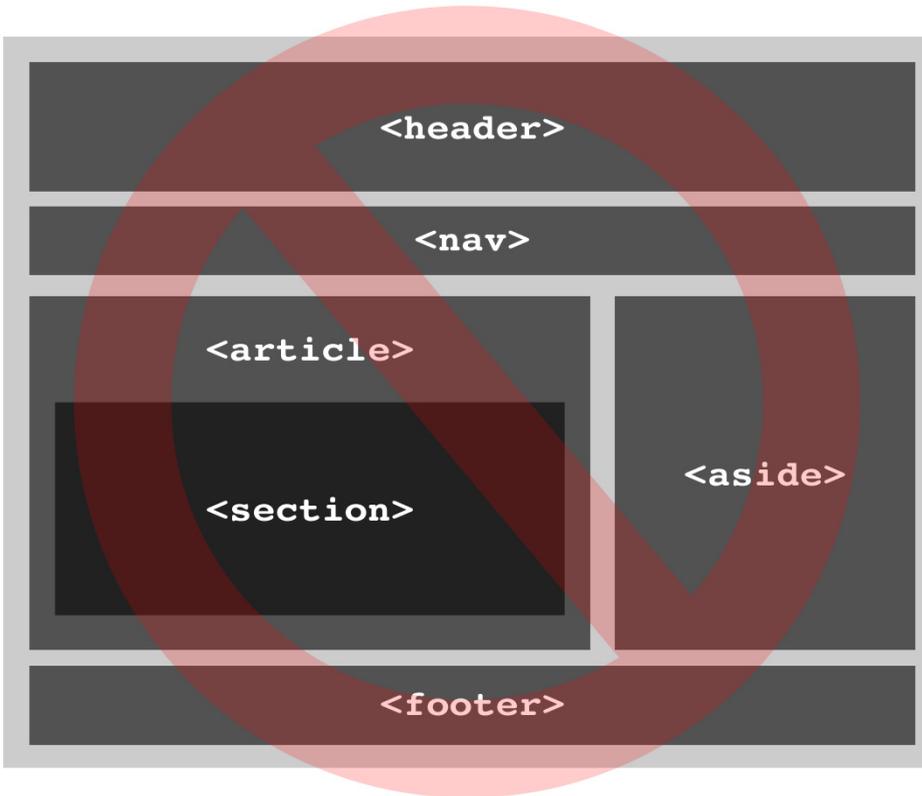


Figure 3.1. This is wrong. Don't do this.

Swapping the new elements in for our old `<div>`s certainly looks easy, right? Nice, clean elements replacing a bunch of random `<div>`s, how lovely!

Unfortunately, there are few problems with this idea:

- **Too few elements.** There aren't enough new elements to do a reasonable 1:1 substitution. Believe me, `<div>`s aren't going anywhere. So if you hear anyone say "Finally, I can get rid of my unsemantic `<div>`s!", you have my permission to pop a cap in their ass with a Nerf gun.
- **Not equal.** While elements are often presented as being equal, they're not. While the "sectioning" elements (`<section>`, `<article>`, `<aside>` and `<nav>`) might work the same, the `<header>` and `<footer>` elements are intended to work *within* the sectioning elements. This can make a *huge*

difference (as we'll see shortly with document outlines), but if you followed most of the discussions about these elements you'd never know.

- **Not a replacement.** When you dig into the HTML5 spec, you discover these tags *as described in the spec* aren't really meant to be 1:1 replacement for existing tags at all. They're actually meant to be used for creating a new form of *document outline*. A document whatnow? We'll explore this next.

These elements have other problems (they don't add anything for semantics or search engines), but we'll talk about them later when we take aim at those two zombie myths that just won't die. We'll also learn what "semantics" actually means in markup, and what search engines really want.

## Outline What Now?

If you try to understand HTML5's new structural elements without understanding *document outlines*, you'll think they're an arbitrary, oddly-named bunch of elements with confusing use cases.

However, once you understand *document outlines*, you'll see they are in fact an arbitrary, oddly-named bunch of elements with confusing use cases that also have an overarching purpose of questionable value.

Granted, this is esoteric stuff. But bear with me and you'll see how HTML5 tries to slip in a radical new way of doing something as fundamental as structuring a web page. This isn't so much paving cowpaths as building a new bovine highway to nowhere.

## What Is An Outline, And Why Should I Care?

An outline is a sort of hierarchical, bullet-point representation of a document.

We actually make an outline whenever we mark up a document and use heading elements. So even if you've never heard of a "document outline", chances are you've already made one. Weird, huh?

The reason we never hear about them is because web designers never have to use them. They're mainly used by blind users as a *primary means of navigation*. When it comes to accessibility, outlines are a big deal. So the best thing we can do to help blind and sight-impaired users navigate documents is provide a good heading structure when using web standards. (We'll explore this a bit more in chapter four).

HTML5 tries to radically change how we make these outlines... and maintain the existing way (well, kind of). This new approach to outlines is the reason the new HTML5 tags exist, and why Hickson and the WHATWG were thinking about adding “section elements” in the first place.

## How We Currently Create Outlines (Even Without Realizing It)

Let's back up a bit and look at our current outlines. In (X)HTML a document's hierarchical structure was dictated through heading levels, using the familiar `<h1>` to `<h6>` tags.

So you might mark up your page like this (as a simplified example), with headings representing the “importance” of each part:

```
<h1>My Sweet Blog</h1>
  <h2>Latest Posts</h2>
    <h3>My Blog Post 1</h3>
    <h3>My Blog Post 2</h3>
    <h3>My Blog Post 3</h3>

  <h3>Blog Sidebar</h3>
    <h4>Blog Archives</h4>
    <h4>Popular posts</h4>
    <h4>Blog roll</h4>

  <h4>Blog Footer</h4>
    <h5>My delicious links</h5>
```

```
<h5>My flickr photos</h5>
<h5>My social networks</h5>
```

The hierarchy or “outline” for the document looks like this:

```
1. My Cool Site
  1. Latest Posts
    1. My Blog Post 1
    2. My Blog Post 2
    3. My Blog Post 3
    4. Blog Sidebar
      1. Blog Archives
      2. Popular posts
      3. Blog roll
      4. Blog Footer
        1. My delicious links
        2. My flickr photos
        3. My social networks
```

Uh-oh. We have a problem. All our lower-level headings are “owned” by the heading above them. “Blog Sidebar” shouldn’t be a heading under “Latest Posts”—it should start a new section.

If we changed the heading level of “Blog Sidebar” to `<h2>` (the same as “Latest Posts”) that would give us:

```
1. My Cool Site
  1. Latest Posts
    1. My Blog Post 1
    2. My Blog Post 2
    3. My Blog Post 3
  2. Blog Sidebar
    1. Blog Archives
    2. Popular posts
```

But now we’re no longer representing the importance of a heading. Instead we’re trying to build a logical structure using a limited set of tags (`<h1>`-`<h6>`), which have a habit of “owning” everything below them—even when they shouldn’t.

Here’s another example. Let’s say we have a page that says:

```
<h2>My HTML5 Book Review</h2>
  <h3>Likes</h3>
    <p>It explained some elements of HTML5 well.</p>
  <h3>Dislikes</h3>
    <p>The author had an annoying habit of writing
silly, self-referential examples.</p>
    <div class="review-body">
      <p>I bought this HTML5 book for the low, low price
of... </p>
    </div>
```

In this document outline, *the entire review* would fall under `<h3>Dislikes</h3>`, because the heading “owns” everything underneath it, even though it should really fall under `<h2>My HTML5 Book Review</h2>`. Usually this *structural* problem goes unnoticed. The *visual* problem of having the review text appear under “Dislikes” however would *not* go unnoticed, so for styling purposes we’d probably introduce a `<div>` so we can visually differentiate between the paragraphs under “Dislikes” and the review body itself.

And indeed that’s often how we structure our documents—we use `<div>`s to break them up into logical sections. But this has no bearing on the document outline in terms of accessibility—the outline is created by headings alone.

As you can see, headings are flawed for creating outlines. People often use heading levels to display different font sizes (with or without CSS), or to indicate arbitrary “importance” rather than structure. And sometimes they just cut-and-paste HTML straight into a new template.

When you consider all that, as well as the limitations of using `<h1>`-`<h6>`, it’s obvious that most web pages don’t have anything *like* a logical outline.

But they *do* have an outline, and using `<h1>-<h6>` *at all* gives blind and sight-impaired users a way to navigate our documents, which research shows is common for people using screen readers. (We'll touch on that research in a moment.) So despite the flaws, for accessibility reasons we need to take headings-for-structure *more* seriously, not less.

(To see the outlines of any site (try your own!), check out the HTML5 Outliner for Google Chrome: <https://chrome.google.com/webstore/detail/afoibpobokebhgfnknfndkgemlggomo>.)

But what if there was a way to create arbitrary outlines *without* relying on headings? As it turns out, people have been thinking about this for years—if not decades.

## “Sectioning” Is An Old Problem

The problem of headings, and how to structure a document, is a long-standing one. XHTML 2.0 proposed a solution in its first draft way back in 2002 (see: <http://www.w3.org/TR/2002/WD-xhtml2-20020805/>), which involved nesting `<section>` tags and using a generic `<h>` element for headings.

This “sectioning” solution in XHTML 2.0 was initially mooted by Tim Berners-Lee as far back as 1991, as Jeremy Keith pointed out (see: <http://adactio.com/journal/1683/>), when Berners-Lee said:

*I would in fact prefer, instead of `<H1>`, `<H2>` etc for headings [those come from the AAP DTD] to have a nestable `<SECTION>..</SECTION>` element, and a generic `<H>..</H>` which at any level within the sections would produce the required level of heading.*

Yep, a good two decades ago.

HTML5 tries to bring this concept of sectioning into mainstream HTML by following a similar path to XHTML 2.0, while also maintaining some backwards compatibility. And the results are, shall we say, *mixed*.

But before we get to HTML5's implementation, let's look at just how important headings are for accessibility.

## If We Care About Blind Users We Should Care About Headings

As we mentioned before, with HTML4 it's headings such as `<h2>Blog Sidebar</h2>` (rather than random `<div>`s such as `<div class="blog-sidebar">Blog Sidebar</div>`) that create document outlines. And for blind users, these headings are important.

How important? In a survey of over 1000 screen reader users (where 80% of the people were blind and 16% were vision impaired):

*The responses to this question provided one of the greatest surprises to us. It is clear that providing a heading structure is important to screen reader users with 76% always or often navigating by headings when they are available. Use of heading navigation increased with screen reader proficiency with them being used always or often by 90.7% of expert users, 79.3% of advanced users, 69.9% of intermediate users, and 55.4% of beginners.*

(You can see the full results here: <http://webaim.org/projects/screenreadersurvey/#headings>.)

Were you aware of that? I wasn't, and I've been using `<h1>-<h6>` willy-nilly for years. I imagine most web designers have some vague idea that `<h1>-<h6>` tags are important, but not how crucial they are for blind users.

So we had an established, straightforward, easy-to-implement way of providing outlines for blind and sight-impaired users. That is, until we hit HTML5.

## HTML5's "Improved" Outlining Was Dead Before It Ever Shipped

We've established what a document outline is (a bullet-point, table of contents-style representation of the page), and we've established how they are currently created (with `<h1>`-`<h6>` elements).

Here, in a nutshell, is how HTML5 proposes to create document outlines:

- Each bullet-point in the outline, or "section", is defined using one of the four "sectioning" elements: `<section>`, `<article>`, `<nav>` and `<aside>`; and not the `<h1>`-`<h6>` elements. The intent here is to solve the limitations of `<h1>`-`<h6>`. (We'll explore each of these new elements in the next chapter.)
- There's no generic `<h>` element, as per XHTML 2.0. But in pure HTML5 it's suggested we could just use `<h1>` everywhere as a generic heading element. In fact, *any* heading element in HTML5 will be treated as a generic heading, with its level determined by how deeply it's nested in sectioning elements.
- But there's no such thing as "pure" HTML5 though, so we need to maintain backwards compatibility. Therefore we should still use `<h1>`-`<h6>` in a logical fashion, which means maintaining *two* somewhat different document outlines in the one document.

That's the general idea. Here's how the spec puts it (<http://www.whatwg.org/specs/web-apps/current-work/multipage/sections.html#headings-and-sections>):

*Sections may contain headings of any rank, but authors are strongly encouraged to either use only h1 elements, or to use elements of the appropriate rank for the section's nesting level.*

Please don't use `<h1>` elements everywhere!

In my view, everyone (Hickson in the spec and in public comments, standards advocates in the community, and designers and authors in general) has made a complete hash of communicating this.

This poor communication has meant that designers and developers have been using these HTML5 elements with no understanding of the outlines they have

been creating. These elements were supposed to bring about better logical document outlines. Instead, given the haphazard way they've been implemented, they've created HTML5-style document outlines that are even *more* broken than the `<h1>`-`<h6>`-based outlines they were intended to replace.

HTML5's version of outlining was effectively dead before anyone understood it, much less implemented it properly.

And here's the irony: this approach, which *may* theoretically deliver accessibility benefits in the future (no one knows when, or even if, screen readers will use these outlines), is destroying the page styles for a small group of IE users *now*. So it's already doing harm, yet has no clear future benefit. (We'll look at this more in the next chapter.)

We'll still explore these new HTML5 elements in chapter four, but mostly so you can understand just how broken they are. (Remember, the cool HTML5 stuff comes in the later chapters.)

## Sneaking In Big Ideas Leads To Dead Ideas

The first problem with this new approach to outlining is the idea that HTML5 is just "paving the cowpaths" and codifying existing practice.

Clearly, introducing a whole new way of structuring documents, however poorly communicated, is not "paving the cowpaths".

You can't then turn around and tell authors and designers, "This is what you've always done!" But Hickson has done just that, saying the new elements are just intended to save common class names. Here are just a couple of examples.

In 2009, Hickson said (<http://lists.w3.org/Archives/Public/public-html/2009Aug/0717.html>):

*They are, more or less, filling the most common requests from Web developers based on what the most common class="" attribute values are. Their main purpose is to simplify authoring and styling.*

And in 2012 (<http://lists.whatwg.org/pipermail/whatwg-whatwg.org/2012-January/034506.html>):

*Mostly these new elements make authoring a bit easier.*

So, if HTML5 is going to introduce a big new idea, it needs to communicate that big new idea. Instead, it appears Hickson doesn't remember, or can't be bothered arguing for, the big idea of sectioning he and the WHATWG added to the spec.

HTML5 advocates (as well as the spec itself) need to communicate the purpose of the new elements properly, or abolish them.

As it is now, they're just inflicting nonsense on the web design community.

Let me give you an example. The spec says `<header>` and `<footer>` elements define areas *within a section*, but do not define sections themselves, and so won't show up in a document outline. This is something most people get wrong, including those teaching HTML5 through books and blogs, whose examples often show `<header>` being on par with `<section>`. The spec also says `<header>` and `<footer>` can be used multiple times per page (once per section, for example), but you would never pick that up from most HTML5 resources out there.

These may seem like pedantic, wonkish points. But they illustrate something very serious—the community is trying to implement HTML5 markup in a way that doesn't have much relation to the actual HTML5 spec. It's a weird in-between state of markup limbo that has inadvertently appeared because that's what everyone *assumed* these elements should be used for.

## We Forked The Spec

In a sense, the community has forked HTML5, as far as markup goes. That's a big problem. There's the “common (but incorrect) understanding” fork of HTML5, and there's the actual HTML5 spec. But following the “common understanding” and replacing visual areas in our templates with elements that “sound about right” has no benefit for anyone. We just create a weird, broken

outline while misusing the new elements. And with so many broken HTML5 outlines there, outlining as a concept is pretty much dead on arrival.

We'll explore each element individually in a moment, but let's stick with the big picture for now.

## How Should We Structure A HTML5 Page?

All this may seem a bit confusing at the moment, so let's take a step back and look at the general rules of structuring a page in HTML5 (such as they are!), as described in the spec:

- We are supposed to use `<section>`, `<article>`, `<nav>` or `<aside>` to create a new *section* in the *outline*. (That is, a new bullet-point in the document outline.) You can see what your outline looks like with the HTML5 Outliner plugin for Chrome: <https://chrome.google.com/webstore/detail/afoibpobokebhgfknfnfdkgemglggomo>. And yes, the terminology here is clumsy—having multiple elements, including `<section>`, create a section in the document outline is quite confusing!
- We use `<header>` or `<footer>` *within each section* to demarcate the header or footer *of that section*. That section can be anything from the root `<body>` section down to an individual comment. (An individual comment is supposed to be an `<article>`, as we'll see in chapter four, which would create a section in the document outline.)
- We use heading elements (`<h1>`-`<h6>`) to give each section a title in the outline and provide backwards compatibility. (As I'm writing this there's no meaningful support for HTML5 outlines anywhere, and doesn't seem to be any on the horizon. So "backwards" compatibility may actually be "compatibility for the foreseeable future".)

You might think you can just replace all your `<div>`s with `<section>`s and create an outline. However, `<section>`s aren't to be used in cases where you only need a styling hook, so in a true HTML5 document you'll still have plenty of `<div>`s. In fact, a "correct" HTML5 document would have:

- a bunch of `<section>`, `<article>`, `<nav>` and `<aside>` tags to create the outline
- a bunch of `<div>`s for styling

- redundant use of `<h1>`-`<h6>` tags to duplicate the outline as best as possible (this is what screen readers will actually use)
- a healthy sprinkling of redundant `<header>` and `<footer>` tags within each section that don't do anything.

Simplifying authoring? With two means of structuring a page, two outlines to maintain, and a bunch of redundant tags to add?

I don't think so. And that's before we even consider styling our headings.

## Styling Headings HTML5 Style Is Kind Of Insane

Let's imagine a pure HTML5 future where we can use `<h1>` everywhere as a generic heading element, as per the spec's suggestion, and we use the new sectioning elements to create the outline. That is, if we use a `<h1>` three sections deep, it's essentially a `<h3>`.

Let's say we want to style this three-sections-deep `<h1>` as though it was a `<h3>`. How are we going to pick it out? And can you imagine picking h1 out everywhere through the cascade to give it a different style for different levels, given four elements can create a section, and can be used in any combination? You wouldn't be able to sleep.

Nicole Sullivan touched on the madness that ensues when you try and style HTML5-style `<h1>` elements through the cascade in the appropriately titled blog post "Don't Style Headings Using HTML5 Sections" (<http://www.stubbornella.org/content/2011/09/06/style-headings-using-html5-sections/>) and gave this simplified example:

```
h1{font-size: 36px}
section h1{font-size: 28px}
section section h1{font-size: 22px}
section section section h1{font-size: 18px}
section section section section h1{font-size: 16px}
section section section section section h1{font-size: 14px}
section section section section section section h1{font-
```

```
size: 13px}
section section section section section section section
h1{font-size: 11px}
```

That, however, is the *vastly simplified* version, as Sullivan points out. The true madness sets in when you have to style all your (say) six-levels-deep headlines which may be nested in any combination of `<section>`, `<article>`, `<aside>`, or `<nav>`. For comedy value, see what such a style sheet would look like here: <https://github.com/cboone/hypsometric-css/blob/master/html5/html5-defaults.css#L426>. It's utterly insane.

The only option then is to fall back on class names for headings, but avoiding class names when authoring is the very “problem” the WHATWG were trying to solve.

And do you think our clients and colleagues who happily create and edit web pages will ever understand the nuances of sectioning their articles correctly? I doubt it.

No wonder people are confused.

Oh, and to top it all off the styling of your `<nav>` (and any other new HTML5 elements) may blow up for ~1% of users. (We'll touch on this again soon.)

This is the way of HTML5. And it's a mess.

Not surprisingly, even the most experienced web authors get bogged down in the HTML5 outline quagmire. Read about Roger Johansson's experience here, for example: [http://www.456bereastreet.com/archive/201103/html5\\_sectioning\\_elements\\_headings\\_and\\_document\\_outlines/](http://www.456bereastreet.com/archive/201103/html5_sectioning_elements_headings_and_document_outlines/).

## This Is Not Inconsequential - People Have To Teach This Stuff

“Ok, maybe the markup wonks got it wrong on this one. Maybe these tags are mostly redundant. So no one uses them, or they don't do it quite correctly. Who cares, Mister Markup Pedant?”

The thing is, introducing these new elements—and concepts such as arbitrary outlining—into the official HTML5 spec means people *actually have to teach this stuff*. (Heck, some designers even teach their kids this stuff—see Cameron Moll’s cool HTML5 whiteboard magnets for example:

<http://cameronmoll.tumblr.com/post/10688505696/html5-whiteboard-magnets>.)

This is bad for web standards. It makes even basic HTML hard to teach, hard to learn, hard to implement, and for what? Structuring a web page should be the least of our worries—not a huge distraction for a generation of students and professionals alike.

(A note to those teaching web standards: if you really hate your students, ask them to explain the difference between `<article>` and `<section>`.)

## Where Does This Leave Us?

Hickson and the WHATWG’s intentions are good. Theoretically, using these tags could improve accessibility even without considering outlining. (For example, screen readers could skip past the `<nav>` tag to get straight to the content.) But the vendors making screen readers have shown very little interest in HTML5 to date. And there is already support for better alternatives, which we’ll look at next.

So we don’t need HTML5’s new elements for accessibility. In fact, we should avoid them for the harm they cause another subset of users.

People will still use these tags, mostly because they want to “do the right thing” in the hope the Standards Fairy will leave small change and/or an Apple product under their pillow. But it’s just a waste of productive time that could be better spent on more important things.

Remember: what ends up in the spec is often just the idea of a few (or even one) interested, smart, ordinary person from (as of writing) 7+ years ago. And chances are even they don’t remember why they wanted it. So I think we’re allowed to disagree about what’s best, and pick and choose what we implement.

But what happens to accessibility? Do we just leave visually impaired users with the status-quo? No, because fortunately there's a better alternative.

## A Sane Approach To Structural Markup For Accessibility

There's a way to add helpers for the blind and visually impaired in our markup *without* wading into the mire of HTML5's new structural elements—ARIA roles.

Actually, it's WAI-ARIA which stands for "Accessibility People Apparently Don't Do Catchy Acronyms". Or, as sticklers for accuracy may tell you, it's "Web Accessibility Initiative - Accessible Rich Internet Applications". (We'll just call it ARIA.)

It's not part of the HTML5 spec. Instead it's a separate (and gigantic) W3C spec that's compatible with HTML5, HTML 4 and XHTML 1.x.

The secret to ARIA is the `role` attribute, which can be added to an element like this:

```
<div role="myariarole">
```

The full ARIA spec is big. Really big. (See it here: <http://www.w3.org/TR/wai-aria/>.) But we'll be looking at a small subset called *landmarks* (see: [http://www.w3.org/TR/wai-aria/roles#landmark\\_roles](http://www.w3.org/TR/wai-aria/roles#landmark_roles) ).

As an example, here are the four main areas of a simple page:

- Header
- Content
- Sidebar
- Footer

And here's how we'd mark it up using ARIA:

```
<div role="banner"></div>
  <div role="main"></div>
```

```
<div role="complementary"></div>  
<div role="contentinfo"></div>
```

Easy.

We'll touch on the roles we can use when we discuss HTML5 elements, and recap in chapter four.

## ARIA Benefits

ARIA roles have several benefits over HTML5 (or previous HTML versions):

- The roles generally reflect how web authors structure pages. (For example, the header, or “banner” is for the stuff at the top of the page—not for every section on the page, as in HTML5.)
- They keep our markup relatively clean, as we can use the `role` attribute as a styling hook for IE7 and above with attribute selectors, such as `div[role="banner"] {border:10px pink;}`. (If you need to support IE6 users, you can also include redundant classes.)
- They work right now in screen readers that support ARIA landmarks, such as JAWS version 10 screen reader, NVDA 2010.1 and VoiceOver on iPhone IOS4+. (See <http://www.paciellogroup.com/blog/2010/10/using-wai-aria-landmark-roles/> for more information.)
- They don't blow up styling for IE6-8 users with JavaScript off like the new HTML5 elements do.

This technique can help blind users now, doesn't hurt web standards, and doesn't require getting your head around a second way of sectioning your document.

We'll look at the appropriate ARIA landmarks to use as we go through the new HTML5 elements in the next chapter.

## Layout Recommendations

Before we finish this chapter, let me recap how I think we should markup pages in the age of HTML5:

- We **should not** use the new tags. (But we'll look at them next, and the ARIA landmarks we should use.)

- We **should** take headings more seriously, given how much blind and sight impaired users rely on them.
- We **should** use ARIA landmarks for accessibility.
- We **should** otherwise use `<div>`s with semantic class names or IDs like we've always done. (If you want to scream "BUT THEY'RE NOT SEMANTIC!", make sure you read chapter five on semantics.)

# THE TRUTH ABOUT HTML5'S STRUCTURAL ELEMENTS

## Okay, Okay. I'll Get To The Tags. But I'm Telling You, It Won't Be Fun

So what have we done so far?

- We've established the broad (and somewhat obscure) concept that HTML5's structural elements are trying to improve—outlining, which is currently done implicitly with heading tags.
- We've established what an outline is for—to help screen readers, who rely heavily on document headings.
- We've also touched on a better way to help blind users get around our pages with ARIA landmarks.

Let's now look at what the HTML5 spec (<http://whatwg.org/html>) says about these new individual elements, starting with...

### **<header>**

Two things you need to know about `<header>`:

1. It doesn't actually do anything.
2. It's intended use isn't quite what you think it is.

The `<header>` element is a good example where a commonly-used term has a new meaning in HTML5, while still being used to “pave the cowpaths”. You probably use `<div id="header">` all the time, so calling it `<header>` will

make it easier to read if nothing else, right? Well, here's one of HTML5's everyone-uses-it-so-let's-change-the-meaning-anyway moments. Here's what the spec says:

*The header element represents a group of introductory or navigational aids.*

*Note: A header element is intended to usually contain the section's heading (an h1–h6 element or an hgroup element), but this is not required. The header element can also be used to wrap a section's table of contents, a search form, or any relevant logos.*

The note is instructive—the intended purpose of `<header>` is to contain a *section's heading*. Remember, sections are created with one of four sectioning elements (article, section, nav, aside) and generate a document outline in HTML5. The `<header>` element is meant to work *within* sectioning elements. It doesn't create a section on its own (despite how it's often visually represented as another sectioning element), and it doesn't add to the document outline.

Think of `<header>` as something that wraps the heading of a section, which can be anything from the topmost section of a document, such as the heading inside the `<body>` tag (i.e. the logo and all that “header” stuff we normally consider a header), right down to the header of a comment.

### Really, It Doesn't Do Anything

All the `<header>` element does is say “this is the header of a given section”. Trouble is, while that's what your markup might say, at the moment none of the browsers or user agents are listening.

And according to Hickson, *they probably never will be* (see Hickson's comment in “Conclusion: R.I.P. HTML5 structural tags” later in this chapter). So this element doesn't do anything now, and probably won't do anything in the future. It's the semantic equivalent of a tree falling in a forest with no one around to hear it.

Given the `<header>` element doesn't modify or add to the document outline, the *actual* heading you see as a bullet point in the document outline (e.g. “My Great

Blog”) is still set by a `<h1>-<h6>` element. The `<header>` tags just wrap those heading elements, along with any other header-y stuff, such as a date. So you could do something like this:

```
<header>
<h1>My blog post</h1>
<p>Published on...</p>
</header>
```

### How Can Screen Readers Use `<header>` When It’s Everywhere?

You may think screen readers could skip the `<header>` element and go straight to the content. But there’s no way we (or the user agents) can be sure the first `<header>` in a document is the main page header. If your mark up is in a non-standard order (the content appears first, followed by the header, footer, and sidebar), the document could have many `<header>`s, none of which we’d call a typical “header”. And so we’re back to square one in dealing with the “overall” header for a page.

### ARIA Alternative: Banner

Fortunately for blind users, there’s an alternative. The ARIA landmark `banner` demarcates the “header” as we currently know it. Here’s how the ARIA spec defines the banner landmark (<http://www.w3.org/TR/wai-aria/roles#banner>):

*A region that contains mostly site-oriented content, rather than page-specific content.*

*Site-oriented content typically includes things such as the logo or identity of the site sponsor, and site-specific search tool. A banner usually appears at the top of the page and typically spans the full width.*

It should appear only once per document so screen readers can jump straight there and be fairly sure what it is—which is exactly what we’re after.

## Recommendation

The `<header>` element is too broad (and too pointless) to be useful. Instead, use the more specific ARIA `role="banner"` on the appropriate element (with a redundant “banner” class for IE6 if need be) for the traditional “header” of a page.

## `<nav>`

Here's the spec:

*The nav element represents a section of a page that links to other pages or to parts within the page: a section with navigation links.*

*Not all groups of links on a page need to be in a nav element—only sections that consist of major navigation blocks are appropriate for the nav element. In particular, it is common for footers to have a short list of links to various pages of a site, such as the terms of service, the home page, and a copyright page. The footer element alone is sufficient for such cases, without a nav element.*

The `<nav>` element *does* create a new section in your outline, and benefits from:

- being somewhat self-explanatory
- having a seemingly useful purpose.

The idea is if we mark up our navigation with the `<nav>` tag, blind people can bypass it and go straight to the content, and jump straight to the navigation links when they want to go somewhere else.

A win for accessibility, right?

## Good Intentions; Accessibility Disaster

Despite these good intentions, doing it for one minority can potentially screw up the navigation for another: IE 6, 7 and 8 users with JavaScript disabled.

Due to the way IE6-8 handles ‘unknown’ elements, these users won’t get any CSS for this element. It could affect one in a hundred users—a higher percentage than people who use screen readers—making the whole idea of using this element somewhat moot for the short to medium term. (This is a problem for all HTML5 elements which we’ll discuss further later in this chapter.)

### ARIA Alternative: Navigation

Fortunately, we can use the ARIA landmark `navigation` instead, by including `role="navigation"` on the appropriate `<div>` (or `<ul>`) to make our navigation more accessible without hurting accessibility for others. The ARIA spec defines the navigation landmark (<http://www.w3.org/TR/wai-aria/roles#navigation>) as:

*A collection of navigational elements (usually links) for navigating the document or related documents.*

### Recommendation

Use `role="navigation"`. Consider the `<nav>` element harmful until only a very small number of IE8 users are left. (This effectively means Windows XP users, as IE8 is the last browser they’ll get. We may be waiting a while.)

### `<section>` and `<article>`

These sound the same (everyone gets them confused), but they have different supposed uses. We’ll look at each one separately first, and then compare the two. Please try and refrain from throwing inanimate objects or small animals while trying to get your head around them.

### `<section>`

Here’s the spec:

*The section element represents a generic section of a document or application. A section, in this context, is a thematic grouping of content, typically with a heading.*

*Examples of sections would be chapters, the various tabbed pages in a tabbed dialog box, or the numbered sections of a thesis. A Web site's home page could be split into sections for an introduction, news items, and contact information.*

*Note: Authors are encouraged to use the article element instead of the section element when it would make sense to syndicate the contents of the element.*

*Note: The section element is not a generic container element. When an element is needed for styling purposes or as a convenience for scripting, authors are encouraged to use the div element instead. A general rule is that the section element is appropriate only if the element's contents would be listed explicitly in the document's outline.*

Okay, let's try and make sense of this. The `<section>` element is supposed to represent a generic section in a document. So if this chapter was a web page, we could break it up into chunks with `<section>` tags. It can also represent different areas of a home page—from news items to contact information. But it shouldn't be used as a generic container element for styling—that requires a `<div>`.

It also shouldn't be used for the page's content area (and neither should `<article>`), but we'll get to that in a moment when we discuss the missing `<content>` element.

## Sections == Outlines

Again, the key to understanding `<section>` is understanding document outlining and the concept of sectioning a document. The spec mentions this (read the last sentence of the second note), but that's not a lot to go on when `<section>` is the main workhorse for creating a document's outline. As a rule

of thumb, as far as creating an outline goes, if it's not an `<article>`, `<nav>`, or `<aside>`, it's probably a `<section>`.

This is also why you shouldn't use `<section>` for generic containers to style. If you just throw them in so you can style an area *without* thinking about your outline, you'll get an illogical, broken outline that defeats the whole point of using `<section>`. It's a common mistake, and shows just how poorly they've been explained in the spec, advocated by the experts, and understood by the community (not that I blame the community).

## Russian Dolls

Don't forget: you can nest sections (whether created by a `<section>`, `<article>`, `<nav>` or `<aside>`). And as we saw in chapter three, in pure HTML5 land this determines the true heading level for a `<h1>`-`<h6>` element—not the level of the heading you use. In HTML5, the user agent (in theory) just sees them all as generic heading elements when inside a section.

So we could use `<h1>`s everywhere, and the user agent would figure out if they were nested as `<h1>` or `<h101>`. However, for screen readers (now, and probably long into the future) we'll need to use `<h1>`-`<h6>` headings appropriately, no matter what flavor of HTML we use.

## Recommendation

If you want to create outlines the HTML5 way, you'll be mostly relying on `<section>`s. It took more than 20 years for the element to make it into the spec (recall Tim Berners-Lee comment in the previous chapter), and it will probably take another 20 years before people understand it correctly.

There's no ARIA equivalent.

## `<article>`

You might think “article” is like “a newspaper article”. Well, shame on you for thinking a new HTML5 element would have an intuitive meaning. Here, it's

more like “an article of clothing”. Yep, another “semantic” term with an unintuitive meaning.

Here’s the spec:

*The article element represents a self-contained composition in a document, page, application, or site and that is, in principle, independently distributable or reusable, e.g. in syndication. This could be a forum post, a magazine or newspaper article, a blog entry, a user-submitted comment, an interactive widget or gadget, or any other independent item of content.*

*When article elements are nested, the inner article elements represent articles that are in principle related to the contents of the outer article. For instance, a blog entry on a site that accepts user-submitted comments could represent the comments as article elements nested within the article element for the blog entry.*

And here’s Hickson on the WHATWG mailing list in early 2012

(<http://lists.whatwg.org/pipermail/whatwg-whatwg.org/2012-January/034506.html>):

*<article> covers a wide range of semantics:*

- forum posts*
- newspaper articles*
- magazine articles*
- books*
- blog posts*
- comment on a forum post*
- comment on a newspaper article*
- comment on a magazine article*
- comment on a blog post*
- an embeddable interactive widget*
- a post with a photograph on a social network*
- a comment on a photograph on a social network*
- a specification*

- *an e-mail*
- *a reply to an e-mail*

In HTML4 a paragraph is a paragraph is a paragraph. In HTML5 an “article” is a forum post is a blog comment is a widget is an actual article. If an element has such broad meaning, how can it be more “semantic”? It’s like deciding to call knives, forks, spoons, plates, and televisions all “forks”.

This isn’t paving the cowpaths.

Again, it’s best understood in terms of outlines. The `<article>` element is for creating a section when you don’t want to use `<section>`, which is usually when you’re wrapping some chunk of content (or “interactive widget”, as the case may be).

The spec talks about how an `<article>` could be syndicated as a self-contained element, but how and why this would ever happen is unclear (use RSS!) It’s a solution looking for a problem.

### Specifications Should Specify

The main problem with `<article>` is it’s open to interpretation (“What does ‘in principle’ mean? “Reusable’?”). Specifications fail when they leave things up to you to work out. The whole point of a specification is to *specify* exactly what you should do. But here it’s open to interpretation, has no clear benefit, and repeats existing functionality (it’s `<section>` with a different name).

### Nesting `<article>` For Articles And Comments

You can also nest `<article>`s within `<article>`s when the content is related. The spec suggests blog comment be wrapped in `<article>` tags, and then nested in an overall `<article>` for the blog entry. This is the “fork” and “spoon” are all “forks” problem. If you have `<article class="post">` and `<article class="comment">` then `article` just becomes a more verbose form of `div`.

Why not just add a `<comment>` element, and at least have basic markup for the standard article-followed-by-comments pattern that just about every blog and publication on the web uses? Wouldn't that be paving the cowpaths? Not according to Ian Hickson, who injects his own idiosyncratic view on the semantics of the issue: that there is no difference between an article and a comment (<http://lists.whatwg.org/pipermail/whatwg-whatwg.org/2012-January/034506.html>):

*I think it's anachronistic to consider that the utterances of the site owner are in some way distinct from the utterances of the site readers. What makes them different?*

*On the contrary, on the Web there is no difference. An article is just a comment that has been hoisted to a more prominent position.*

The irony of defining at least one of the differences between these “utterances” and then declaring there is no difference was apparently lost on Hickson.

Of course, this also flies in the face of the oft-stated goal that these elements are mostly to help authors maintain their documents. If ever there was a case where a pattern of markup has emerged—an article, followed by comments—this is it. Yet Hickson, acting as both player and referee, wont budge, asserts his own peculiar philosophical view that all “comment” is equal, and that's that. We're left with nesting `<article>`s in HTML5, and content on the web, it seems, will be `<article>`s all the way down.

### Search Engines Don't Need `<article>`

Some may think `<article>` could help search engines, but they don't need a tag to know where your content is. Their entire existence relies on them being able to find your content without that sort of help. Even if `<article>` was used widely, how would they know through mark up alone whether `<article>` meant blog post, forum post, interactive widget, comment or whatever? It's too broad to be useful for SEO even if they *did* care what tags you used (which for the most part they don't).

(We'll discuss this more in chapter six when we look at HTML5 and SEO.)

## **<article> Is Not For The “Main” Part**

And `<article>` isn't for denoting the content area of the page. There's no actual element for this—it's the missing `<content>` problem we'll get to in a moment.

I'm yet to confirm whether any illicit substances were involved in the defining of this element.

### **Recommendation**

Should you use this? I wouldn't. Instead I'd put it in the guilty-of-being-useless-until-proven-otherwise category. If a pragmatic benefit appears, go nuts. Until then, pass.

Like `<section>`, there's no ARIA equivalent.

## **So What's The Difference Between `<article>` And `<section>`?**

Some things you need to know:

- Articles can be nested within articles.
- Articles can be broken up by section.
- A section can be broken up into articles, which can in turn have individual sections.
- People are terrible at using markup consistently.

Guess what? Apart from the handful of markup uber-nerds, anyone who actually uses these elements (and I'd be surprised if anyone does) will just create a huge mess. But hey, maybe I'm wrong.

Personally, I'd rather take a rusty potato peeler to my pinky than debate the virtues of `<section>` vs `<article>`. The very fact there's a debate at all demonstrates a failure in the spec. If you have to debate an element when implementing it, you lose.

Yet thousands of words about this fine distinction have been spilled on blog posts (for example, Bruce Lawson's take: <http://www.brucelawson.co.uk/2010/html5-articles-and-sections-whats-the-difference/>) and comment threads.

One would have hoped the absurdity of the situation would have dawned on the community when we're left making well intentioned, but otherwise ridiculous *flow charts* just to decide which HTML5 element to use (see: <http://html5doctor.com/downloads/h5d-sectioning-flowchart.png>). Alas, it appears it has not. And all this because a WHATWG member or three decided to throw these additional `<section>` flavors into the Web Applications 1.0 spec in 2004.

Ok, we've ripped off the band-aid and survived the most painful part of HTML5's new elements. But there's still some sticky residue that will sting a bit coming off, so let's look at the final sectioning element.

## `<aside>`

Question: What do you call a pull quote, parenthetical information, and a sidebar?

If you said, “a pull quote, parenthetical information, and a sidebar” you lose—buh-bow. You call them an “aside”. Obvious, huh? Here's the spec:

*The aside element represents a section of a page that consists of content that is tangentially related to the content around the aside element, and which could be considered separate from that content. Such sections are often represented as sidebars in printed typography.*

*The element can be used for typographical effects like pull quotes or sidebars, for advertising, for groups of nav elements, and for other content that is considered separate from the main content of the page.*

It's hardly obvious why a pull quote, advertising, or a blog's navigation elements (complete with `<section>`s for a blog roll and archive—that's the use cases in the spec) should be called the same thing. But in the bizarro-world of HTML5 structural semantics they are, and you can happily ignore them.

## Aside Creates An Outline Section In Weird Places

Aside *does* create a section in a document outline, which makes it all the more strange given its broad use cases. (Why should a pull quote be its own section?)

If it was just for a sidebar, or even called `<sidebar>` (as it originally was), that might make sense, but it's not, and it doesn't.

### ARIA Alternative: Complementary

The ARIA landmark alternative in this case is `complementary`, which the ARIA spec describes as (<http://www.w3.org/TR/wai-aria/roles#complementary>):

*A supporting section of the document, designed to be complementary to the main content at a similar level in the DOM hierarchy, but remains meaningful when separated from the main content.*

*There are various types of content that would appropriately have this role. For example, in the case of a portal, this may include but not be limited to show times, current weather, related articles, or stocks to watch. The complementary role indicates that contained content is relevant to the main content. If the complementary content is completely separable main content, it may be appropriate to use a more general role.*

### Recommendation

Use `role="complementary"` on the appropriate `<div>` (or other element) for the sidebar in your template (and anywhere else that's appropriate).

### `<footer>`

Remember how `<header>` seemed obvious but wasn't? Well, it's the same deal with `<footer>`. It sounds like it should just be the main footer, but in fact it can mean the footer for any given section. Here's the HTML5 spec:

*The footer element represents a footer for its nearest ancestor sectioning content or sectioning root element. A footer typically contains information about its section such as who wrote it, links to related documents, copyright data, and the like.*

*Contact information for the author or editor of a section belongs in an address element, possibly itself inside a footer.*

*Footers don't necessarily have to appear at the end of a section, though they usually do.*

Footers, like headers, are something you use within a section—they don't create a section in their own right. Again, this is confusing because in the examples in HTML5 books and blog posts the footer appears to be a distinct *visual* section in a template of equal importance to an `<article>` or `<aside>`, when in fact it doesn't appear in the document outline at all. It simply describes part of the parent section, wherever (and however frequently) that section may occur.

### Footer Doesn't Do Anything Either

Like `<header>`, `<footer>` doesn't actually do anything. In fact, `<footer>` doesn't even have to go at the end of an element. The spec gives an example in the `<article>` section where the metadata of a comment (comment author and timestamp) is wrapped in a `<footer>` and placed at the—wait for it—*top* of the comment. Apparently the element wasn't confusing enough.

Again, no cowpath paving here. I certainly didn't hear a thousand web designers crying out for a `<footer>` element for every section on a page. Did you?

### Fat Footer? Good Luck!

What if you have a trendy “fat footer” with a bunch of links and other information in it? Well, if it's a whole chunk of content it should be a section, and so you'd use one of the four sectioning elements—`<nav>`, `<section>` or even `<aside>`.

But, just to keep things interesting, you *can* wrap that sectioning element in `<footer>` tags. So `<footer>` can both demarcate content *within* a section in a document outline, *and* wrap other sections in the document outline, and yet still not create a section in its own right.

Clear as mud, huh? I don't understand why it's that way, and I doubt most web designers will either. I've seen Japanese game shows that make more sense than these elements.

### Can I 'ave A Footer, Guv'na?

Proposals for a `<footer>` element have been around for a long time. When it was discussed in 2002—about a decade ago—in regards to XHTML 2.0 on the W3C public mailing list, people soon began criticizing it (see: <http://lists.w3.org/Archives/Public/www-html/2002Aug/0257.html>) saying it wouldn't provide any tangible benefit. A decade later that criticism still seems valid.

### ARIA Alternative: Contentinfo

Again, as far as accessibility goes, ARIA to the rescue. Its `contentinfo` landmark reflects the content of traditional footers (i.e. the bit at the bottom of the page with a few links and some fine print) without a presentational name. The ARIA spec describes it as (<http://www.w3.org/TR/wai-aria/roles#contentinfo>):

*A large perceivable region that contains information about the parent document.*

*Examples of information included in this region of the page are copyrights and links to privacy statements.*

Unlike `<footer>`, it should only be used once per document.

### Recommendation

We can use `<div role="contentinfo">` once for the page footer. We can also keep using `<div id="footer">` (or whatever you like) with `role="navigation"` if there's substantial navigation elements there.

## Where Is My `<content>` Element?

So far, we've covered four sectioning elements (`<nav>`, `<section>`, `<article>` and `<aside>`) and two elements that define areas within those sections (`<header>` and `<footer>`).

These elements mostly define sections in a document outline, which is something I doubt many web designers will bother with. But let's say you decide to use them anyway. You start marking up your overall template, creating sections with the four sectioning tags, and dropping in `<header>`s and `<footer>`s where appropriate.

But wait, something's missing. What tag do you use for the actual content area?

A `<div>`.

No, I'm not being a smart-ass. There really is no `<content>` or `<main>` tag.

Why not? Ian Hickson thinks it's redundant. By definition, anything in a perfectly marked up HTML5 page that *isn't* in a `<header>`, `<nav>`, `<aside>` or `<footer>` (and presumably top-level `<section>`s, but let's not go there) is the `<content>` for the page.

Here's Hickson explaining the decision on Twitter (<http://twitter.com/#!/Hixie/status/8922228541>):

*[A]nything that \_isn't\_ in `<header>`, `<footer>`, `<aside>`, or `<nav>` is by definition `<content>`, so we don't need an explicit element.*

While the logic is cute, telling web designers to mark up different parts of a page with explicit (and often redundant) tags, but not to worry about the actual content (the main part of the page) because its markup is implied, is absurd.

When something is implied like this, you have to be “in the know”.

Unfortunately most web designers won't be “in the know” about this special rule, and so will unknowingly misuse `<section>` or `<article>` to fill in the perceived gap.

Bruce Lawson explained what he'd seen on <http://html5doctor.com> (<http://lists.w3.org/Archives/Public/public-html/2009Aug/1366.html>):

*Regardless of purity, authors \*want\* to have a `<main>` tag.*

*HTML5 doctor gets lots of emails about how to mark up main content, and it's arguably the biggest misuse of `<section>` as people are using `<section>` to do it. (HTML5 doctor itself does it like this; we know, and plan to change it)*

*Perhaps we've all be brainwashed by those evil Web Standards People, but it doesn't seem right that you mark up the peripheral stuff with their own elements, but the main content -the purpose of the page- merely gets a measley meaningless generic `<div>`.*

(Jeremy Keith also chimed in with his support of a `<main>` element:

<http://lists.w3.org/Archives/Public/public-html/2009Aug/1397.html>.)

Assuming web authors will correctly use the other HTML5 structural elements *and* understand their “main” content should use just a `<div>` instead of an explicit tag is terrible. It would be far better to just add the tag to the spec so user agents could jump straight to it (assuming they decided to supported it).

But despite this objection being raised several times over the years of HTML5's development, Hickson disagreed, and that's that. NO SOUP TAG FOR YOU.

### **ARIA Alternative: Main**

Fortunately there's yet another ARIA landmark we can use instead, which is simply `main`. It does what we want it to do, and serves the people we're trying to help—blind users.

Here's how the ARIA spec defines `main` (<http://www.w3.org/TR/wai-aria/roles#main>):

*This marks the content that is directly related to or expands upon the central topic of the document. The main role is a non-obtrusive alternative for "skip to main content" links, where the navigation*

*option to go to the main content (or other landmarks) is provided by the user agent through a dialog or by assistive technologies.*

## Recommendation

Use `<div role="main">` once for the main content area of your page.

## Other ARIA Landmarks

Here are a few more ARIA landmarks that may prove useful when marking up your pages:

- `application` for software widgets on a page.
- `form` for forms, except search, which gets...
- `search` for the search form on the page.

And they go along with the others we've looked at:

- `banner` for the overall header.
- `navigation` for, you guessed it, navigation.
- `complementary` for side bars.
- `contentinfo` for the footer.
- `main` for the main content area.

(Note that `banner`, `main` and `contentinfo` should only be used once per document.)

For more about these landmarks, see Steve Faulkner's excellent comparison to HTML5 elements: <http://www.paciellogroup.com/blog/2010/10/using-wai-aria-landmark-roles/>.

For more on ARIA in general, see Mozilla's ARIA documentation: <https://developer.mozilla.org/en/aria>.

## A Funny Thing Happened... Graceful Degradation Died And JavaScript Became Mandatory

In the last couple of chapters I've made several references to the harm these elements can inflict on another small subset of users. For some users, these tags will blow up your page in IE6, 7 or 8 if they have JavaScript off (which, due to personal choice and over-zealous security concerns, is more common than you may think).

The problem is how IE 6-8 handle “generic” elements. To IE6-8 a generic tag is *anything* it doesn't recognize, whether it's a completely made up element like `<mymadeuptag>` or HTML5's `<nav>`. Given that IE6-8 doesn't recognize generic elements at all, we can't style them without using a sprinkling of JavaScript to tell IE6-8 they exist.

This clever JavaScript workaround was discovered by Sjoerd Visscher and popularized by Remy Sharp in 2009 (see: <http://html5doctor.com/how-to-get-html5-working-in-ie-and-firefox-2/>). Now we can use a tiny bit of JavaScript to tell IE6-8 that particular elements *do* exist, we can happily style away... or can we?

No, we can't.

But could we?

No.

(For the record, the JavaScript in question is available here: <http://code.google.com/p/html5shiv/>. You also need to tell other browsers to treat the new elements as block-level elements with CSS, as we touched on in chapter two.)

Now we hit the big snag: the IE 6-8 users with JavaScript off who don't get the HTML5 fix.

Hang on. Does anyone actually turn JavaScript off these days? Is there any data available that could give us some insight here?

Yes, there is.

## Yahoo's JavaScript Research

In 2010, Yahoo published the results of research they did into this very question—how many visitors *do* have JavaScript disabled? It turns out that 2.06% of visitors hitting Yahoo's US websites (which includes significant non-US traffic) had JavaScript disabled, as did 1.29% in the UK, 1.46% in France, and 1.28% in Spain. (Brazil was an outlier with just 0.26%).

So, unless you're designing for Brazilian users, you still need to consider users who've disabled JavaScript.

You can read about the research here: <http://developer.yahoo.com/blogs/ydn/posts/2010/10/how-many-users-have-javascript-disabled/> and the methodology here: <http://developer.yahoo.com/blogs/ydn/posts/2010/10/followup-how-many-users-have-javascript-disabled/> on Yahoo's YDN blog.

Let's assume the high end of 2% for traffic to the major US sites, and that IE 6-8 users make up about 50% of your audience. That's at least one in every hundred users with IE6-8 and JavaScript off. Even if it were one in a *thousand*, that's still at least one person every day for a moderately busy site.

So what happens to these visitors if you use HTML5 structural tags?

If you use `<section>` tags (or any other tags) around existing markup and *don't* style them, nothing. You can still create outlines that way if you really want, as long as you remember that blind users still rely on headings to get around your page.

But if you use `<header>`, `<footer>`, `<nav>`, `<aside>` or `<article>` (or `<section>` for that matter) where you *are* likely to style them, Bad Things™ happen.

Specifically, users with JavaScript disabled will see your site styled normally, *except* for the parts of the design that use the HTML5 elements. When those elements wrap our navigation, header(s), sidebar(s), or article(s), we have problems. The elements—unlike the rest of the page—won't have our CSS styling applied, and so may break in some pretty serious ways.

## Here's What Happens...

For example, here's the main navigation of one of my client's sites, firstly with `<div id="nav">` wrapping a normal `<ul>` list of links:



Figure 4.1. The site's navigation with JavaScript disabled and no HTML5 elements.

And here it is with a `<nav>` element instead in IE7 with JavaScript disabled:

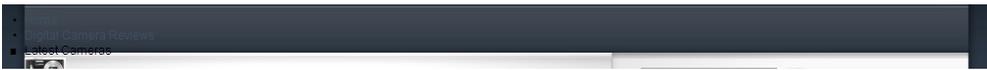


Figure 4.2. The same navigation using HTML5's `<nav>` element with JavaScript disabled.

Not cool, right? And that's just one element. Imagine if the header and sidebar broke down in similar ways, while the body, wrapper and content areas were all still being styled. Not pretty.

## What To Do? Oh, XP....

The safest thing would be to not give IE6-8 users with JavaScript disabled styles at all (which you could do with conditional statements and writing the style element in with JavaScript). But I don't advocate this—you're still unnecessarily ruining their experience.

What about IE9 and beyond? Thankfully IE9 recognizes HTML5 elements, and has improved how it handles generic elements.

Maybe IE9 will replace IE6, 7 and 8 sooner rather than later, and we can forget about these non-JS users. Ah, what a world that would be. Unfortunately, Windows XP users will never get IE9—IE8 is the end of the road for them. As

long as XP is around we'll have IE8 visitors whose experience is unnecessarily ruined because of these elements.

And no-one should have to put up with that.

Elements such as `<nav>` are supposed to (theoretically) help the visually impaired. But when we *do* use `<nav>` (for example) as intended in the spec, and as taught by HTML5 advocates as a replacement for existing structural markup, we unnecessarily hurt another minority in a very real way.

## Uh... Web Design Community, What Happened?

One of the most unfortunate parts of the HTML5 hype in web design circles is how quickly we've forgotten the tradition of graceful degradation. Sure, we don't want to be limited to the lowest common denominator. But IE8 is a long way from that, and there's a big difference between giving those with JavaScript disabled something simpler, and giving them something that's just plain broken.

Let me make it clear I don't have a problem if *your specific website* requires JavaScript (and will break in varying ways for IE6-8 with JavaScript off), and that's the informed choice you've made. But I do have a problem with this issue being almost entirely ignored in the web design community, and the unnecessary pain unwitting designers and developers are inflicting on a small group of users.

This issue alone sinks the usefulness of these new elements, and so I recommend sticking to `<div>`s for structure and ARIA roles for accessibility. ARIA roles can at least help solve accessibility problems for one minority without blowing the site up for another.

As Hickson himself says (<http://lists.whatwg.org/pipermail/whatwg-whatwg.org/2012-January/034506.html>):

*Naturally, if you are happy with `<div class="...">` for everything, you are welcome to continue doing that.*

## Conclusion: R.I.P. HTML5 Structural Elements

Perhaps the strangest thing about all this is I'm not sure Ian Hickson—the spec editor—is clear about what these new elements are for, or what other experts think they'll be used for.

In an exchange with Opera's Bruce Lawson (co-author of *Introducing HTML5* (New Riders, 2010)) on the WHATWG list in 2009 (<http://lists.whatwg.org/htdig.cgi/whatwg-whatwg.org/2009-March/018888.html>), Lawson said (emphasis added):

*After all, I know of no user agents that can use time, section, footer, datagrid etc **but we mostly expect there to be soon.***

And here's what Hickson said in reply:

*I don't. Most of the new elements are just meant to make styling easier, so that we don't have to use classes.*

That's Hickson's offhand rationale—not having to use classes—which we touched on earlier. But it's interesting he has no expectation user agents will do anything with them. Has he given up on sectioning and outlining? Isn't that why these elements were added in the first place (remember they predated any research), and isn't that how Hickson described them in the specification he edits?

These elements are a lost cause. If people *do* just use these elements instead of classes (which they already are), they won't be considering the outline they'll be creating. They may assume `<header>` creates a section (it doesn't), or use `<aside>` for a pull quote (and in doing so create a new section), which will mean a lot of broken HTML5 outlines.

So user agents (screen readers in particular) will have very little incentive to ever use HTML5 outlines—the main point of these elements—as a means of navigation, given they'll be (and already are) an utter mess.

It's sad, isn't it? Tim Berners-Lee's 1991 wish for sectioning finally makes it into a shipping HTML spec (after the aborted XHTML 2.0 project), only to be misunderstood and abused two decades later.

And that's *before* we take into account the ambiguities in the spec, the fact they blow up styling in IE6-8, and how much they complicate *marking up a basic HTML page*, of all things.

Here's John Allsopp (author of *Developing with Web Standards*, New Riders, 2009), commenting on the new elements on Jeffrey Zeldman's blog (<http://www.zeldman.com/2009/07/13/html-5-nav-ambiguity-resolved/#comment-44699>):

*As I waded deeply through the specification, narrowing my focus largely to the new "semantic" elements like section, article, header, footer, and so on I discovered many ambiguities, poorly explained features, and baffling containment rules [as well as] near byzantine complexities.*

*[...] There's recently been much jubilation that XHTML2 is dead. But many of the legitimate criticisms of XHTML2 can also be laid at the door of HTML 5. And many others, particular to the specification itself can also be leveled. It's time to address these.*

We have the solution for marking up our pages:

- use `<div>`s with classes, (and/or a single `id` that occurs once on the page if you want)
- use appropriate heading elements
- add ARIA landmarks as appropriate and you're done.

That's how easy it is to describe structural markup now.

That simplicity has been lost in HTML5.

That's the bad news. The good news is, all is not lost. We can take this opportunity to try using headings in a better structural way to make blind and sight-impaired users lives a bit easier (and start using ARIA landmarks for the

same reason). And we can be happy knowing that our existing HTML structural markup techniques will serve us just fine for many years to come.

Let me finish by saying that despite my criticism of this particular part of the HTML5 spec (and the web design community's advocacy for it), if it wasn't for Ian Hickson and the WHATWG there wouldn't *be* a "HTML5" as we know it today—we could still be waiting for the W3C to get their act together. It's not my aim to bite the hand that feeds—just nibble on its fingers a little bit.

Next, let's talk about the 'S' word: semantics.



# THE TRUTH ABOUT HTML5 MICRO-SEMANTICS AND SCHEMA.ORG

### Let's Talk About Semantics, Baby.

A common claim made about the new HTML5 structural elements is that they are “more semantic”.

In my view, the new elements are “more semantic” in the same way fruit-flavored candy bars are “more nutritious”—not at all.

Nevertheless, the question of semantics in HTML5 gives us an excellent excuse to take a quick trip through the big picture of “semantic” markup. We’ll look at where semantic markup came from; what semantic markup promised to deliver but never quite did; and we’ll finish with a quick look at something you can use right now—new schemas put forward by the major search engine companies (Google, Microsoft and Yahoo) that will hopefully improve the display of your search results.

By the end of this chapter your markup nerd-dar will be so finely tuned you’ll be able to separate the markup poseurs using “semantic” as a mere buzzword from the hardcore markup wonks who are still waiting for the Semantic Web to arrive, any day now...

## Semantics In A Nutshell

When it comes to the web, there are actually two kinds of “semantics”—the nitty gritty markup of a given web page, and the so-called Semantic Web. Let’s start with the semantic markup we practice every day as web designers.

“Semantic markup” was one of the cornerstones of the web standards movement. In 2003 Jeffrey Zeldman, perhaps the best-known advocate for semantic markup and web standards, wrote on his blog (<http://www.zeldman.com/daily/0303a.shtml>):

*CSS combined with lean semantic markup makes sites faster, more portable, and more accessible. The combination helps sites work in more existing environments and is the best hope of preparing them for environments that have not yet been developed.*

This was a major change in both theory and practice for web designers. We’d keep all the styling information about a page in a separate CSS file, and describe the content with “lean, semantic markup”, as Zeldman put it.

Here’s a (slightly reworked) example of semantic markup Zeldman used in a 2002 Digital Web article ([http://www.digital-web.com/articles/999\\_of\\_websites\\_are\\_obsolete/](http://www.digital-web.com/articles/999_of_websites_are_obsolete/)). Firstly, Zeldman borrowed some “unsemantic” markup from an ecommerce site to show what we were moving away from (try not to shudder when you read it):

```
<td width="100%"><font face="verdana,helvetica,arial"
size="+1" color="#CCCC66"><span class="header"><b>Join now!
</b></span></font></td>
```

And then, with CSS handling the styling, the markup simply became:

```
<h2>Join now!</h2>
```

And lo and behold, there it was: lean, semantic markup that we pretty much take for granted now. It was a big, and extremely worthwhile, shift in practice.

But what makes this example “semantic” and not the first one? “Semantic” is just a fancy way of saying “meaningful”, and by using heading tags (<h2>) it

now means something to browsers (and screen readers)—“This is a heading”. Screen readers can (and do) use these headings to navigate around a document, and browsers can give these elements default styling (e.g. making it a block-level element).

It also makes it easy for us humans to read as well. When we scan the markup there’s no doubt about what this text is—it’s a heading. Simple, right?

This highlights the two key groups that matter in “semantic markup”: humans and machines (browsers, screen readers, search engines, etc.) It should be both “human readable” and “machine readable”—Semantic Markup 101.

“Machine readable” semantic markup has other benefits. Search engines can scan, index and search our content in a way that’s much harder (if not impossible) with Flash sites, or websites consisting purely of images (as print designers are occasionally wont to churn out).

That said, Google doesn’t care much about what markup you use.

## These Problems Have Been Solved

Here’s the thing: the solution to these problems has been around for more than a decade, no matter what flavor of HTML you are using. Search engines can index our content, screen readers can understand it, and our lean, semantic mark up makes it easy to read and maintain.

Then the pedants took over.

The people in web design circles began to think “Well, if *semantic* is good, then *more semantic* must be better, right?”

Not really.

Beyond the point of human readability, and basic machine readability, “more semantic” doesn’t mean anything (irony ahoy!). But this hasn’t stopped people debating which elements are more semantic or more appropriate, which nine times out of ten is about as useful as debating whether it’s “splade” (or more correctly, “splayd” for you sticklers out there) or “spork”. (Splade, obviously.)

## There's No Such Thing As "More" Semantic

I humbly propose that the unqualified use of "more semantic" be banned from web design discussions about HTML elements post haste.

Whenever you hear someone going on about something being "more semantic", ask them this simple question.

"For who?"

If all they can come back with is "But it's ... MORE SEMANTIC!", they're just making a vague claim about nothing. But if they say something like "More semantic for screen readers", *that's* a valid claim we can evaluate.

Do screen readers *really* do anything different for these "more semantic" elements? Are they supported at all? Or do they cause bugs like the HTML5 elements did when they were first used? (See: <http://www.accessibleculture.org/blog/2010/11/html5-plus-aria-sanity-check/>)

(Remember: due to the no-JS IE6-8 issues, using HTML5 elements for accessibility is like punching one kid to entertain another one.)

Likewise, if they say "But it's more semantic *for search engines*", we can evaluate that specific claim. What does Google's developer guidelines say? What does the SEO community think? And so on.

But please, no more unqualified claims of "But it's *more semantic*" when discussing HTML5. These dubious assumptions have been attaching themselves like barnacles to the good ship Web Standards for years, and it's time we revved up the high-pressure hose and cleaned them off (assuming that's how barnacles are, in fact, removed).

Ok, mini-rant over. The human readability and basic machine readability problems have been solved, this is where we're at, and we may hope that HTML5 will take us forward. But before we get to HTML5's approach, let's talk about the Big Idea™ behind semantic markup.

## Big Ideas In Semantic Markup—The Semantic Web

What if we could take the “machine readable” part of semantic markup further? What if the machines (and browsers in particular) could read our markup and know not just what content appeared, but what given blocks of content actually meant?

That’s the big idea behind semantic markup. If we can describe the content of our pages *accurately* and *specifically*, then machines can do cool stuff with the data.

This is (or perhaps was) partly the idea behind the “Semantic Web”—a big, broad concept that would be driven by the XML-ified web. (Read more about it here: [http://en.wikipedia.org/wiki/Semantic\\_Web](http://en.wikipedia.org/wiki/Semantic_Web).) The web would be a perfectly described library of documents, marked up in excruciating detail with XML. An XML-based future was something many influential people believed in. In fact, in the earlier markup example from 2002 and the use of `<h2>`, Zeldman described web standards as a way we can “transition from HTML, the language of the Web's past, to XML, the language of its future”.

However, as we saw in chapter one, the move to XML died, and with it the dream of a true “Semantic Web”. Instead, the web became a wonderful platform for applications, went social, and kept on being the web we know and love. But it wasn’t the capital-S Semantic web people had hoped for.

We need to keep this history in mind when people talk about “semantic” elements in any situation, whether it’s HTML5 or whatever future HTML evolves. What kind of “semantics” are they referring to—basic human and machine readable semantics we all use every day, or the dead-end dream of the XML-powered Semantic Web?

## Semantics: Not Dead Yet, Or, Google & Co Drop A Micro-Semantic Bombshell

There’s actually a third option that sits between the lean, semantic markup we use now and the pie-in-the-sky Semantic Web—microdata (and microformats), which adds a layer of metadata to our markup.

(A variety of approaches compete here, particularly microformats, microdata, and RDFa. But I'll just be referring to the overall concept as “micro-semantic”, which is also known as “structured data”.)

With micro-semantic, we simply embed semantic data into our existing HTML document. Let's look at how micro-semantic could help daily life on the web.

## Ecommerce With Real (Micro) Semantics

Let's use online shopping as an example. Here, truly semantic markup could *theoretically* help desktop browsers (i.e. all of us), the visually impaired using screen readers, and search engines.

**Desktop browsers:** Let's say we're shopping online for a new TV, and doing our research by visiting a bunch of websites to compare features and prices for specific models. In most cases (well, if you're as obsessive about research as I am) this means copying and pasting the relevant information from each site into a separate document—which is both tedious and prone to error.

Now, imagine if these e-commerce sites all marked up their pages with a `<productdetail>` tag and nested `<price>` and `<specs>` tags. Our browser could easily find the product detail part of the page, and with a single click we could add the price and specs into a comparison shopping list. With specific, meaningful tags, your browser—a machine—can find, compile and sort certain information for you very quickly. After all, that's what they do best.

**Screen readers:** It could also help the blind or vision-impaired. Imagine a blind person doing the same research for a new sound system. If the e-commerce pages were marked up with these `<price>` and `<specs>` tags, their screen reader could theoretically read out *just* the price and product specs. They could then save those details to their comparison shopping list, and move on. But until it happens they have to try and navigate around highly complex pages by having headings and content read out to them.

**Search engines:** With the prices and specs marked up correctly, Google, Bing and other search engines could display the price of a product fairly reliably in

their search results and improve the whole search experience. (This is actually possible right now, which we'll get to.)

They're just a few examples of what's possible when we have truly semantic markup. Machines—browsers, screen readers, and search engines—can easily pick out useful information, and do cool things with it (such as create a comparison shopping list).

Problem is, to use different tags to describe this data the HTML spec would need a squillion different tags. Every kind of content—from poems to products to policy documents—would need its own tags so the machines knew what the content was. The list of HTML tags would literally be a small dictionary, or rather a very *large* dictionary as more and more tags were added to the spec. Authors writing about HTML would quite likely lose their minds.

The good news is we *can* mark up our content and make this comparison shopping possible (especially the search engine example) without needing any more HTML tags. We simply annotate our existing HTML with attributes and values that machines can read. (I'll talk more about this soon.)

Adding a handful of new elements HTML5-style, however, is not a path to “more semantic” documents. They don't help machines do much with the data, and our markup becomes more cluttered—hardly a way to make it more readable.

Instead, we need a new mechanism to describe this data. Hopefully that's where HTML5 will lead us.

## Can The Real Semantics Please Stand Up?

I know what you're thinking. “If only we had a way of adding tags that didn't pollute the entire spec. Some sort of eXtensible Markup Language.” But as we saw in chapter one, we tried that and it failed.

Clearly we need a way to extend HTML that doesn't involve a) adding a dictionary's worth of elements to the spec, or b) trying to XML-ify the web.

There is a third option, and a bunch of people have been working on various solutions for quite a few years.

Here's the idea in a nutshell: just attach attributes with values from an agreed bunch of terms to our existing HTML. Here's an example (I've made up the attribute and value):

```
<div class="myclass" semanticdata="mysemanticvalue"> ...
content ...</div>
```

As you can see, it's pretty simple. But it's worth teasing out the terminology, as the different terms and implementations can make a simple idea seem far more complex than it actually is.

We need to distinguish between several pieces of the micro-semantic pie:

- **The infrastructure:** We can take different technical approaches when *adding* semantic data to a document (i.e. what HTML infrastructure we use). It boils down to which attributes we use—the existing `class` attribute (microformats), the new HTML5 attributes such as `itemprop` (microdata), or attributes such as `property` and `content` (RDFa). Not surprisingly, people interested in the nitty-gritty get all worked up over which is best. But it's *what* we say, not *how* we say it, that's much more interesting.
- **The vocabularies:** What we say—the *kind* of data that we stick in these attributes—is where the rubber hits the road. And we need to work together to make it work. The people implementing the data (web designers), and the companies that might do something with it (e.g. search engines) need to agree on a stable set of terms—a vocabulary—to describe a review, person, or event so everyone's on the same page.
- **The concept:** And then there are the communities built around different ways to implement this infrastructure and vocabularies, and do cool things with it.

One group who have been doing cool things with micro-semantic data is the microformats community. They have an active community (<http://microformats.org/>), a microformats way to use HTML as infrastructure (the `class` attribute), and specific microformat vocabularies. These are the various parts of the micro-semantic pie, and demonstrate how communities have been able to come together to do semantics in a meaningful way on the web.

You may have heard of, and perhaps implemented microformats in the past. Unfortunately, as I write, its future has been more or less killed off by the search giants who have proposed a new way forward for micro-semantic or “structured data”.

## Why Should We Care About Micro-Semantics?

As I write this, Google, Microsoft and Yahoo! have launched what may be the biggest effort to get real semantics into HTML documents in the history of the web.

And how did they launch it? With a blog post and a website that had all the pizzazz of a “My First HTML Page” template knocked up during a hurried lunch break. And they also managed to single-handedly piss off everyone already invested in the process who’ve been evangelizing micro-semantic for years. Not a good start.

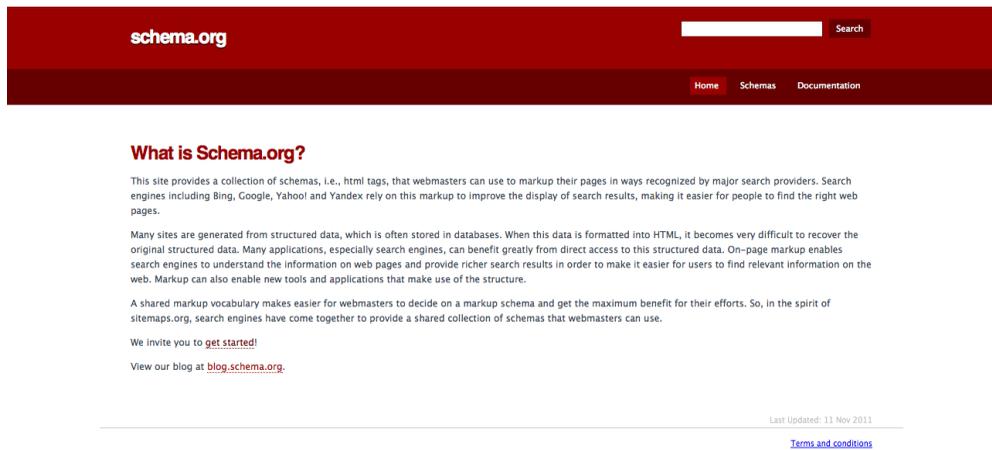


Figure 5.1. Schema.org. Who said semantics weren't sexy? Oh... everyone. Right..

## Schema.org—The Future Of Semantics?

In mid-2011 a handful of engineers across Google, Microsoft and Yahoo! decided they didn't like the current, community-driven approaches, and announced they were picking HTML5's microdata as the winning *infrastructure* (i.e. the HTML attributes we should use to add micro-semantic data). And so they released Schema.org (<http://schema.org/>)—a list of vocabularies, or “schemas”, that the major search engines would use to display richer search results.

In this way, all three parts of the micro-semantic pie were changed. The infrastructure (HTML5's microdata), the vocabularies (Schema.org) and the drivers (corporations, not communities) were all new.

(You can read Google's announcement here:

<http://googlewebmastercentral.blogspot.com/2011/06/introducing-schemaorg-search-engines.html> and Microsoft's announcement here: [http://www.bing.com/community/site\\_blogs/b/search/archive/2011/06/02/bing-google-and-yahoo-unite-to-build-the-web-of-objects.aspx](http://www.bing.com/community/site_blogs/b/search/archive/2011/06/02/bing-google-and-yahoo-unite-to-build-the-web-of-objects.aspx).)

Here's an example of a richer search result:

**Apple iPhone 4 Review - Smartphones & PDA Phones**   
 ★★★★★ Review by Joseph Hanlon - 30 Jul 2010  
 30 Jul 2010 ... Some will still consider Apple's iPhone the benchmark for today's smartphones, but its revised design raises issues we haven't seen in a ...  
[www.cnet.com.au](#) > [Mobile Phones](#) > [Reviews](#) - [Cached](#) - [Similar](#)

*Figure 5.2. Google “iPhone review” and you’ll get results similar to this one. Note how much metadata is included—rating, reviewer, date, and breadcrumbs are all present here.*

## Couldn't We Do This Before?

This is similar to the “Rich Snippets” micro-semantic initiative Google launched in 2009, which you may have heard about (or even implemented). But Rich Snippets only supported a handful of existing vocabularies, and let authors

choose between microdata, microformats, and RDFa. (Plus it was only supported by Google.)

Now we have one “approved” infrastructure for implementation (microdata), one set of vocabularies at a central location, and a big reason for implementing them: support in Google, Bing and Yahoo!.

That’s a big deal.

(Keep in mind this is purely for search result *display*, not search *ranking*. It’s important our clients know the difference.)

What’s remarkable isn’t the search giants choosing one infrastructure, but rather the 300-odd vocabularies that will potentially define semantics on the web for years to come. And it was all done behind closed doors with no standards process (or community involvement) whatsoever.

## The Semantic Web We’ve Been Waiting For?

Make no mistake—this is the biggest, actually-supported thing to happen for semantics on the web since, well, pretty much forever.

Way back in chapter one we looked at how XML was supposed to transform semantics on the web but didn’t. (It was just Architecture Astronauts at work.) We’ve also looked at how HTML5 adds a few semantic elements that are either harmful or add up to very little. (Adding more elements to HTML proper isn’t a solution for semantics.)

This approach of micro-semantics promises a middle way interested communities have been exploring for some time. Let’s run through the existing approaches before we look at the schema.org launch (and everything that was so horribly wrong with it).

## Microformats

The microformats community has been developing and advocating micro-semantics with reasonable success for years, after kicking off in 2004 (see:

<http://microformats.org/wiki/history-of-microformats>). From <http://microformats.org/about>:

*Designed for humans first and machines second, microformats are a set of simple, open data formats built upon existing and widely adopted standards. Instead of throwing away what works today, microformats intend to solve simpler problems first by adapting to current behaviors and usage patterns (e.g. XHTML, blogging).*

For example, in February 2011, all of Facebook’s events were published using microformats (see: <http://microformats.org/2011/02/17/facebook-adds-hcalendar-hcard>). And with the appropriate browser extension (such as the Google Calendar extension for Chrome) a button would appear next to an event, which you could click to add the details to your calendar. Pretty neat, eh?

How did Tantek Çelik, one of the founders of microformats.org, react to the Google and Microsoft schema.org announcement (<http://twitter.com/#!/t/status/77083481494142976>)?

*#schemaorg spits in the eyes of every person and company that worked on open vocabularies like vCard, iCalendar, etc.*

Ouch.

## RDFa

Microformats was a simple, straightforward, limited-by-design approach to micro-semantic.

RDFa (or Resource Description Framework—in—attributes) was the W3C’s much more complex (but more flexible) approach to machine readable data that’s been kicking around since 1997 as just “RDF”. (RDFa was started in 2004.) It never really captured developer interest in any significant way, but it’s still hanging around.

As debate raged about the Schema.org announcement mid-June, Mark Pilgrim quipped (<http://twitter.com/#!/diveintomark/status/80980932957450240>—link now 404s; this was before Pilgrim’s internet disappearing act):

*The W3C: failing to make RDF palatable since 1997*

Zing.

But there have been some interesting real world uses, such as the GoodRelations vocabulary for ecommerce (<http://www.heppnetz.de/projects/goodrelations/>) that could drive the ecommerce example we looked at earlier.

Web designers generally prefer the simplicity of microformats to the flexibility and complexity of RDFa. Nevertheless, a community interested in micro-semanticism had grown around RDFa.

How did Manu Sporny, the current Chair of the W3C’s RDF Working Group, react to Google and Microsoft’s schema.org announcement? In “The False Choice of Schema.org” (<http://manu.sporny.org/2011/false-choice/>):

*Schema.org is the work of only a handful of people under the guise of three very large companies. It is not the community of thousands of Web Developers that RDFa and Microformats relied upon to build truly open standards. This is not how we do things on the Web.*

Yikes.

## Microdata

Finally we have microdata, the new format used in schema.org.

Nothing compels web authors to add esoteric metadata to their pages like several competing, slightly different metadata formats. So Ian Hickson, the HTML5 editor, decided microformats was too cold, RDFa was too hot, and so invented a third approach—microdata—that he felt was *just right* (so to speak). (Here’s Hickson’s lengthy WHATWG post introducing the feature: <http://lists.whatwg.org/htdig.cgi/whatwg-whatwg.org/2009-May/019681.html>.)

Note that microdata, as far as the HTML5 spec is concerned, is about providing the *infrastructure* (with new, valid attributes) for adding micro-semantics. It doesn't specify what those vocabularies should be, or who should invent or maintain them. It is completely separate from the actual vocabularies on schema.org (for example).

And this is the format that won, in a blessed-by-the-tech-giants sense.

(For a lengthier discussion of the various formats, and the implications of schema.org, see Henri Sivonen's excellent "Schema.org and Pre-Existing Communities" <http://hsivonen.iki.fi/schema-org-and-communities/>.)

## Microdata And Schema.org

Now Google, Microsoft and Yahoo! are pushing not only a single format (microdata), but also a single set of vocabularies too for real semantics on the web.

Everything has a specific vocabulary (or "schema"): books, movies, events, organizations, places, people, restaurants, products, reviews, you name it. (See the full list here: <http://schema.org/docs/full.html>.)

There are even schemas for identifying parts of web pages themselves, including the header, footer, sidebar, and navigation. I guess ARIA, HTML5 and so on weren't enough.

If this takes off, and that's a *big* "if", it will be a huge revolution in how we mark up our pages—bigger than XHTML, HTML5, and whatever flavor of HTML comes next.

Has the "Semantic Web" finally arrived?

## How Not To Launch A New Initiative

*"schema.org... there's just nothing quite like throwing away years of vocabulary/ontology work"*

—Jay Myers, 3 June 2011, <http://twitter.com/#!/jaymyers/status/76344419867037696>

Well, not if the tepid launch of this new initiative is anything to go by. It was pretty much a textbook case of what *not* to do.

Here are a few things they could have handled *slightly* better:

- **Consultation:** The Schema.org announcement came from nowhere—no consultation with the community, no heads up, just a desire to “get something out there”.
- **Outreach:** It’s generally not a great idea piss off the people who’ve spent years advocating something similar to what you’re launching. Instead of getting the microformats and/or RDFa communities on board (or at least encouraging a migration path), Google, Microsoft and Yahoo! completely ignored them. And that made them very unhappy campers.
- **Human face:** Schema.org launched as an utterly generic site with almost no human aspect, just a “feedback” button. Who edits it? Who thought it up? Who do we talk to? What’s the process? *Is* there a process? It’s a complete mystery as far as the website goes. Indeed, the FAQ asks “Who is managing schema.org on an ongoing basis?” and answers “Google, Bing and Yahoo! are managing schema.org on an ongoing basis.” Right, well I guess we can always just contact Google, Bing and Yahoo! then. (To be fair, they eventually got a Schema.org blog up here: <http://blog.schema.org/>.)
- **Newbie friendly:** For most website designers, micro-semantics is a pretty “out there” concept. And while Schema.org does have a “Getting started” guide (<http://schema.org/docs/gs.html>), it needs a much friendlier explanation of the how, why and what of micro-semantics (including schema examples) if they want anyone besides the in-the-know uber nerds to pick it up.



For example (slightly abridged):

**Ivan Herman:** *Schema.org is out there, ... how do you envisage the process for the future whereby schema.org might be a place where new vocabs are developed. I [sic] place to make it a more open social process?*

**Kavi Goel:** *I don't have a great answer right now. I don't think any one company wants to own this in its entirety. By going with 3, we showed we [Google] weren't just doing it. [...]*

*Then it leaves the question of where is the completely open discussion... We don't have an answer yet, but this is important. We'll need to sort out the stuff that's out there.*

**Kevin Marks:** *Ours [microformats] has an edit button, yours has a feedback button. The CORE of microformats is we reach agreement. YOU said "we did it in a closed room". You haven't shown your work, your evidence, how others can get involved. This is the most worrying thing.*

**Kavi Goel:** *That's a totally valid point. Microformats did a great job creating an open community.*

*There's no good answer for why we didn't do that.*

*Coming to microformats with a whole bunch of new things could have been an option. We did want to get something out there.*

Earlier in the discussion, Goel said:

*The achievement was to get something out there. We know it's not perfect. We can make it better. We hope this can be a step toward great adoption.*

Here's hoping. The rush to "get something out there" seems to have done more harm than good at this stage, but they can redeem themselves. We now have one format, and one set of vocabularies to use for micro-semantics on the web. If

Google (and/or Microsoft) actually throws some resources at it, and someone at either company actually takes ownership of the project, it could be a very big deal indeed.

To the credit of those involved in Schema.org, consultation is finally taking place, and interested parties are discussing a way forward. See, for example: “Schema.org Workshop—A Path Forward” [http://semanticweb.com/schema-org-workshop-a-path-forward\\_b23387](http://semanticweb.com/schema-org-workshop-a-path-forward_b23387).

Also see the sporadically-updated Schema.org blog for further outreach efforts: <http://blog.schema.org/>.

## Wrapping Up: Semantics And HTML

The waves from the Schema.org announcement are still rippling out across the web as I write. But even so, we can still say a few things about semantics, HTML, and what we should do:

- **The semantic problem:** True semantics that describe *meaning*, and not just *structure*, happens in a layer *on top of* HTML. This seems to be the solution to the long standing problem of semantics on the web. XML won’t bring us true semantics, nor will more HTML tags. It’s a layer of micro-semantics on top of our existing HTML that will.
- **Microformats and RDFa are probably dead-ends:** Microformats has done really well over the years, and I love its simple format. But the decision by Google and Microsoft makes these formats look like dead ends, and the micro-semantics ecosystem (including browser add-ons and validators) will presumably move to microdata and Schema.org vocabularies. Of course, Schema.org could flop hopelessly too, and the microformats community (for one) could keep plugging away. (Google’s not *dropping* any support, in any case.)
- **Get involved:** It’s worth reading up on and experimenting with microformats tools that already exist (such as browser extensions and bookmarklets) to get a taste for what’s possible with micro-semantics. But the fact Schema.org appears to be the future means we as a community need to study the various schemas and provide feedback.
- **Questions remain:** Many questions about the process (will there be one?) and future of Schema.org remain unanswered—questions even the instigators can’t answer, as Kavi Goel demonstrated. And there are bigger

questions about its widespread use. Will mainstream adoption lead to attempts to spam search engines? (People will certainly try.) Will it all turn to “metacrap” (<http://www.well.com/~doctorow/metacrap.htm>)? We will see.

- **It’s ready to go:** Google and Microsoft’s “Better to ask for forgiveness than permission” attitude with Schema.org means the standards process won’t be going on for years—it’s good to go right now. And if it is widely adopted, our online shopping example may eventually become a reality. For now, here’s the February 2012 announcement of using Schema.org micro-semantic to describe videos, which is “now the recommended way to describe videos on the web”: <http://googlewebmastercentral.blogspot.com/2012/02/using-schemaorg-markup-for-videos.html>.
- **It’s being used right now:** Companies like **eBay**, **IMDB**, **Rotten Tomatoes**, and others have implemented Schema.org’s semantics and are benefiting from improved display of their search engine results right now, as this article demonstrates: <http://www.seomoz.org/blog/schema-examples>.

Ultimately, Schema.org is a case of glass half-full, glass half-empty. We now have a well-supported, standard set of semantic schemas we can easily add to any HTML structure. And if we search with Google, Bing or Yahoo! we can get tangible results. The chicken-and-egg problem of adding semantic data has been solved, the format has been chosen, and the schemas have been released.

But rushing the launch (which was underwhelming, to say the least), abandoning any standards process whatsoever for the vocabularies, and trampling years of existing work is a heavy price to pay.



# THE TRUTH ABOUT HTML5 AND SEO

### The Markup And SEO Myth

One odd myth that keeps perpetuating in books and blogs is that new (or even old) HTML elements will help with search engine optimization (SEO). Let's put this one to bed right now (with phasers set to "rant"), and consider the broader question of markup and SEO.

SEO is all about ranking well for given search terms in a search engine's index (usually Google, but it could be Bing, or even China's Baidu). For example, I might hope this book ranks well for "HTML5 book" (hey, I can dream). But this ranking has very little to do with markup—semantic or otherwise. (I'm simplifying greatly here - industry websites like <http://searchengineland.com> and <http://searchenginewatch.com> demonstrate how vast the world of modern SEO is, but we'll keep it simple here and focus on the issue of markup and rankings.)

### SEO In The Dark Ages

A million years ago, search engines ranked your page by just looking at the content of your page—what keywords were used, where they were used, and how frequently they appeared. To this day people still think they need the "keywords" meta tag (`<meta name="keywords" content="redundant, page, keywords">`) because Google uses it in its rankings.

It doesn't. Google has been ignoring it for years because people abused the hell out of it. (See: <http://googlewebmastercentral.blogspot.com/2009/09/google-does-not-use-keywords-meta-tag.html>.) Today the keywords meta tag only serves as a useful indicator to see which SEO "experts" are still stuck in the 90s.

## Stuff Your Keywords

Remember the old joke "How many SEO experts does it take to change a light bulb, lightbulb, light, bulb, lamp, lighting, lightswitch, switch, energy"?

Google doesn't care as much as we'd like to think about the keywords on our sites, because people lie, fudge, and cheat when it comes to the data they control. (For a broader look at this phenomenon with metadata, see "metacrap": <http://en.wikipedia.org/wiki/Metacrap>.)

Google's breakthrough innovation was to look at what *other* web sites said about a given site by looking at the quality, quantity, and content of links pointing back to it. That is, "off-page" factors determine how your web site ranks for a given search term far more than the "on-page" metadata. (It's also why comment spam took off—people tried to game Google's rankings by posting spammy links everywhere, thinking more links would result in better rankings.)

## HTML And SEO

The new HTML5 elements are a form of metadata—data about data—and Google couldn't really care less whether you use an `<article>`, `<footer>`, `<div>`, or even a `<table>` to structure your page.

(No really, Google doesn't care if you use tables for layout. Which isn't to say you should, but it gives you an idea of what they're dealing with. See Matt Cutts' explanation here: [http://www.youtube.com/watch?v=fL\\_GZwoC2uQ](http://www.youtube.com/watch?v=fL_GZwoC2uQ).)

Google has to make sense of the web as it is—the glorious, hideous mess of invalid tag soup—to get the best information possible back to its users. To put it another way, the burden is on Google (and Bing, and Baidu) to rank the web,

rather than on every single web author (including those wielding Front Page) to provide perfectly marked up pages to be blessed by the ranking fairy.

The 0.000001% of pages that use valid HTML5 are largely immaterial. Google won't crawl your page and say "Wow, there's an `<article>` tag, I'm definitely going to rank this page higher than the next guy!" (But if you believe that, please see my next book "30 Incredible HTML5 SEO Secrets Guaranteed To SuperCharge Your Rankings!")

## But What If It Helped... Somehow?

SEO can be a ferociously competitive business, where the basic strategy is usually getting more and better links than your competitors (see this article and discussion on a modern approach to "link building": <http://www.seomoz.org/blog/strategic-link-building-why-you-dont-need-to-outrun-lions>).

If we're serious about SEO, this is the sort of thing we need to focus on. Yet despite the huge significance of a site's link profile, some people still insist that maybe, just maybe, using the new HTML5 elements will "help" search engines parse your content, and therefore improve your search ranking. But that's like saying Kobe Bryant could be "helped" with some suggestions on how to play basketball. You can assume they've got it sussed over at the Googleplex by now.

Again, this doesn't mean we should write sloppy markup—far from it. We still have to maintain it, after all. It's just that new, fancy markup and search engine rankings have little to do with each other.

And we should provide the search engines with more metadata for search result *display* when they ask for it. (Schema.org may prove helpful in this regard, as we saw in the previous chapter.) Metadata for search result *display* is far less likely to be abused and therefore it should remain useful. We just shouldn't believe more metadata means better rankings.

But let's say I'm wrong and Google did look favorably on HTML5 tags. Guess what would happen? People would abuse the hell out of them to try and get an advantage, and Google would remove whatever benefit it gave those tags to begin with. Such is the arms race between the army of engineers working on

good quality search results, and the many site owners and SEO consultants dedicated to manipulating them.

(That “manipulation” isn’t always bad, mind you—it is both possible and preferable to add value by producing great content that people link to, raising your search rankings the good ol’ fashion way, as described by Matt Gemmell in “SEO for Non-dicks” <http://mattgimmell.com/2011/09/20/seo-for-non-dicks/>.)

## Zombie Myths Must Die... Eventually

When web standards first took off, we *could* argue that web standards were good for SEO, insofar as they were better than a Flash page that made it impossible for a search bot to “see” the links and content buried in a Flash file. Similarly, print designers who exported their entire sites as images weren’t doing themselves any SEO favors, as text in images is nigh on useless. But as far as Google is concerned, if enough good links come back to a site then that’s highly relevant to its ranking—even if it’s a pure Flash site.

This myth about SEO and markup probably grew from these beginnings. If basic text and markup helps search engines, then *better* markup must help them *even more*, right?

No. It’s time to put this myth to bed: HTML5 does not help SEO. The difference in ranking between you and the next guy is not because of the tags you use on a given page. HTML5 brings a lot of interesting things to the table (particularly as we’ll see in the coming chapters), but improved search engine rankings is not one of them. HTML5 for SEO is about as effective as homeopathy.

# THE TRUTH ABOUT HTML5'S OTHER NEW ELEMENTS

## Other HTML5 Elements: The Good, The Bad, And The Kind Of Nutty

We've covered structural markup in some depth, now let's get down to the nitty gritty. HTML5 redefines several inline and block-level elements, and introduces a few more. We'll run through some of the changes and additions, and then consider the broader philosophy behind these elements.

### Be Bold Or Die Trying

Let's start our look at these elements with something as seemingly innocuous as the `<b>`, `<i>`, `<strong>` and `<em>` tags, and what's changed in HTML5.

Only in web standards land can we turn something as straightforward as **bold** and *italic* into a complicated mix of dogmatism, high-level theory, and broken pragmatic reality.

When web standards took off, we all endeavored to separate *presentation* from *content*. No longer would font tags and tables clutter our markup. Instead we'd style our pages with CSS, and describe our content in a meaningful (i.e. semantic) way with appropriate tags.

This put the poor old `<b>` and `<i>` tags in a tough spot. They were ostensibly *presentational*—they described how text should look, not what it meant—and

we were running away from presentational tags as fast as our fingers would carry us. So we all embraced the `<em>` tag instead of `<i>` (for “emphasis”), and the `<strong>` tag instead of `<b>` (for “strong emphasis”). These new tags now described the meaning of the text—it was *emphasized*, and how “emphasized” text looked (or sounded) was (theoretically) up to the browser or screen reader. We could then use `<b>` and `<i>` for purely stylistic reasons, and `<strong>` and `<em>` for semantic purposes—a subtle difference, but a difference nevertheless.

I embraced the change (you may have too), thinking that mattered. But it didn’t. Yes it helped draw a line in the sand between presentational and semantic markup, but this was splitting rather narrow hairs. There was no pragmatic benefit. For example:

- **We just swapped one for the other:** Given `<strong>` still bolded text, and `<em>` still italicized text, we all just swapped `<b>` for `<strong>` and `<i>` for `<em>`, and that was that. The difference between what was “emphasized”, and what was just bold or italicized styling *without* any particular emphasis, was lost, given we kept using them for presentational purposes anyway. WYSIWYG editors were particularly guilty of this. The difference was just too subtle.
- **Screen readers ignored them altogether:** The main benefit of these “semantic” tags (supposedly) was that screen readers could read the text with “emphasis” or “strong emphasis”. In fact, screen readers, by and large, ignore them altogether. (See: <http://www.paciellogroup.com/blog/2008/02/screen-readers-lack-emphasis/> for further discussion.)
- **Search engines don’t care:** Google treats `<strong>` and `<b>`, and `<em>` and `<i>` *exactly* the same. (See Matt Cutts video here: <http://video.google.com/videoplay?docid=-1756437348670651505#>.)

So for all the dogmatism about these elements, the reality is pretty simple—use whatever you want. The humans that read it won’t care, and the machines that read it (screen readers, search engines) don’t care either.

But where do these elements fit in HTML5?

I guess if you’re writing a spec you have to *try* to make some sense of how these elements are used, with some emphasis (pun intended) on how they *should* be used. Here’s what the spec says (emphasis added):

`<i>`—The `i` element represents a span of text in an *alternate voice or mood*, or otherwise offset from the normal prose.

`<em>`—The `em` element represents *stress emphasis* of its contents.

`<b>`—The `b` element represents a span of text to be *stylistically offset* from the normal prose without conveying any extra importance.

`<strong>`—The `strong` element represents *strong importance* for its contents.

HTML5doctor.com has an entire article on how this might work in theory (see: <http://html5doctor.com/i-b-em-strong-element/>), but it's really pure fiction. If you think people will actually mark up their documents in this way, I have 15 billion web pages I'd like to show you. As Ian Hickson himself likes to say (<http://www.webstandards.org/2009/05/13/interview-with-ian-hickson-editor-of-the-html-5-specification/>):

*[I]f they [browser vendors] don't implement it, the spec is nothing but a work of fiction. [...] I don't want to be writing fiction.*

If the HTML5 spec documented actual behavior (i.e., “paving cowpaths”), the spec would just say `<b>` and `<strong>` make text bold; `<i>` and `<em>` make text italic; and screen readers tend to ignore them altogether. That's the reality. Everything else is fiction.

This may seem like small fry, but we've touched on a bigger philosophical question: how much of marking up a document in HTML is word processor-like formatting, and how much is marking up the *meaning* of the text? For most web authors—usually our clients using the Content Management Systems we set up for them—it's about word processor-like formatting, *and that's ok*. We'll return to this shortly.

## Wrap Your Anchor Around This, And Other Bits And Pieces

Let's do a quick roundup of some other features and elements available in HTML5

### Wrap Anchors Around Block-Level Elements

We can now do things like wrap a link around a `<h1>` heading and paragraph, which could be useful for items such as blog posts. We need to set the wrapping `<a>` element to `display:block`; (see: <http://mattwilcox.net/sandbox/html5-block-anchor/test.html>) or there could be unexpected behavior.

(Some people have reported problems in Firefox 3.5 (see the discussion in this article [http://www.smashingmagazine.com/2009/08/04/designing-a-html-5-layout-from-scratch/#highlighter\\_741571](http://www.smashingmagazine.com/2009/08/04/designing-a-html-5-layout-from-scratch/#highlighter_741571)), so test thoroughly when wrapping links around block-level elements.)

### `<mark>`

There's a new `<mark>` element we can use to highlight text (with appropriate CSS), instead of, say, `<span class="mark">keyword</span>`. This could highlight search keywords in search results, for example.

### `<figure>` and `<figcaption>`

The `<figure>` and `<figcaption>` elements let us to mark up a photo, chart, table, code snippet, or any other self-contained content that's referenced from "the main flow of the document", as the spec says. So, we might have:

```
<figure>  <figcaption>Yup, this is my photo.</figcaption> </figure>
```

(See the spec for more examples: <http://www.whatwg.org/specs/web-apps/current-work/multipage/grouping-content.html#the-figure-element>.)

These elements may be mildly helpful for accessibility (i.e. screen readers could

read out the figure and its caption), but it's a complex issue. See this extensive write up by Steve Faulkner for more: <http://www.paciellogroup.com/blog/2011/08/html5-accessibility-chops-the-figure-and-figcaption-elements/>.

These elements also suffer from the same IE6-8 no-JS styling problem we discussed earlier.

## <time>

The new <time> element was included mostly for microformats (well before Schema.org was born) but should be useful for future micro-semantic initiatives. Beyond that, <time> is deceptively complex. It's the drama queen of HTML5 elements, and if <time> were a TV show it would be *Bold and the Beautiful*.

In 2011 alone it was killed off by Ian Hickson, then half-revived in the W3C HTML5 spec, and then re-added by Hickson in an improved way to the HTML5-but-we-just-call-it-HTML WHATWG spec. Bruce Lawson blogged about <time>'s removal and reappearance here: <http://www.brucelawson.co.uk/2011/goodbye-html5-time-hello-data/> and here: <http://www.brucelawson.co.uk/2011/the-return-of-time/>.

And it has been subject to a great deal of debate on the WHATWG mailing list before all the 2011 drama (Ian Hickson summed up one debate in 2009 here: <http://lists.whatwg.org/htdig.cgi/whatwg-whatwg.org/2009-March/018888.html>).

It's worth pondering how the HTML5 editor can arbitrarily kill off an element on a whim, add a new one (<data>) and then re-invent the previously-dead element in the face of a backlash, in what is supposed to be a specification that browser makers can implement.

(If you're a sucker for punishment, or are coming off some kind of Charlie Sheen-esque bender and *really* need some sleep, there's also a 8,000+ word WHATWG wiki entry on <time> here: <http://wiki.whatwg.org/wiki/Time>.)

So how do we use this new, risen-from-the-grave version of <time>?

In its current incarnation, the <time> element allows a variety of strings such as a year string (2011), a month string (2011-11), a date string (2011-11-12), a time

string (14:54) with or without seconds and microseconds, combinations of date and time strings (2011-11-12T14:54:39.92922), and more complex strings with time-zone offsets (2011-11-12T06:54:39.92922-08:00).

For example, you could use it like this:

```
<p>The Y2K bug destroyed civilization on
<time>2000-01-01</time>.</p>
```

This is more liberal than the original incarnation of the `<time>` element, and for a full list of valid strings see the spec: <http://www.whatwg.org/specs/web-apps/current-work/#the-time-element>.

The `<time>` element also allows a machine-readable datetime value that can be stuck in the `datetime` attribute, with something more human-friendly in the `<time>` tags (or indeed, nothing at all), such as:

```
<p>The Y2K bug destroyed civilization at the <time
datetime="2010-01-01"> beginning of this year</time>.</p>
```

This is handy for micro-semantics, such as Schema.org microdata.

You can also add a boolean `pubdate` attribute to indicate when an `<article>` (or the overall document if it's not within an `<article>`) was published:

```
<p>This Y2k article published on <time pubdate
datetime="2000-01-01T01:42">Dec 31, 1999</time>.</p>
```

(Remember a *boolean* attribute simply means “yes” by including it, i.e. “this is the publication date”, or “no” by excluding it—it doesn’t accept values.)

### **<details> and <summary>**

The new `<details>` element functions as a show/hide box without having to use JavaScript. It has a boolean attribute (i.e. standalone, with no value) of `open`, which tells the browser to display the box as open by default. But if the attribute is absent, it will be collapsed, with the `<summary>` element describing what appears when collapsed.

Here's an example:

```
<details>
  <summary>Show/hide me</summary>
  <p>You can see this when expanded</p>
</details>
```

This would give the following result:

▶ Show/hide me

---

▼ Show/hide me  
You can see this when expanded

*Figure 7.1. The <details> element closed (above) and open (below).*

The spec suggests it could be used in complex forms (and uses OSX's file info window as an example) where you want show or hide certain settings or form inputs. Browser vendors are still working out how they should style this by default. Currently, only Chrome supports it.

This a strange addition to the spec, and one of the WHATWG's curious little innovations. Common patterns of JavaScript- or CSS-powered behavior have become quite prevalent in recent years (think tabs, drop down menus, pop overs, lightboxes and so on), and yet there's no desire to have that functionality replicated in pure HTML. A show/hide triangle control was, however, deemed worthy of being included in the spec. Such are the little mysteries of the WHATWG's HTML5.

## <small>

Some existing HTML4 elements have also been redefined.

For example, the `<small>` element now means “fine print” not “visually small”. I find the idea of redefining an element this late in the game weird, but there you go.

### `<address>`

I didn't even know there *was* an `<address>` element. It's a block level element that, in HTML5, is for contact information for a given section (e.g. an `<article>`, perhaps in the `<article>`'s `<footer>`), or the document itself. The spec says it's explicitly *not* for arbitrary postal addresses, which should just be in `<p>` tags. If someone from the WHATWG finds out you've used it for an arbitrary postal address, expect to have a finger shaken *very* firmly in your direction.

### `<cite>`

In HTML5 `<cite>` has been redefined to *exclude* the previously-acceptable use of citing people's names. It's now only for works. This really annoyed Jeremy Keith, who wrote about it on 24 Ways (see: <http://24ways.org/2009/incite-a-riot>). Again, it's weird the HTML editor can just redefine elements on a whim. It raises the question of whether we should bother with these elements for “inline semantics” at all, which brings us to...

## Should We Even Use These Obscure Little Tags?

If and when new *functionality* arrives in browsers or other agents (and not just the bowels of the HTML5 specification), sure, some of these elements may prove handy from time to time.

But let's step back and consider the bigger picture of the purely “semantic” text-level elements. We've already touched on the question of simple word processor style formatting versus marking up meaning when discussing `<b>` vs `<strong>` and `<i>` vs `<em>`. Now let's consider the `<address>` element for example. In November 2009, Jack Osborne wrote on HTML5doctor.com (<http://html5doctor.com/the-address-element/>):

*The address element has been around since the HTML3 spec was drafted in 1995, and it continues to survive in the latest drafts of HTML5. But nearly fifteen years after its creation, it's still causing confusion among developers. So how should we be using address in our documents?*

Perhaps, after 15 years, it's time for a rethink. What's our aim here? Are we going to give it another 15 years? After 30 years, will the web finally be using `<address>` correctly? And if it is, so what?

15 years ago we may have assumed that 'one day' someone will do something useful with our carefully marked up pages. We now know better. It's time to re-evaluate. We've spent 15 years experimenting with HTML to see what works in terms of semantics and functionality. It's time to take stock of the results.

If HTML5 was truly paving cowpaths here, it would open up the definition (instead of tightening it) for elements like `<address>` and `<cite>`, or better still make them obsolete altogether. We don't need them. They don't do anything. Micro-semantics *on top of HTML* make them obsolete. The search engines have demonstrated through Schema.org (and earlier initiatives such as Rich Snippets) that they want micro-semantics, not redefined HTML elements. Authors have little use for them. So why keep them?

This is the truth we need to acknowledge when it comes to these finer aspects of markup. HTML for documents has proven to be pretty lousy for anything but basic semantics which are explicitly tied to formatting (header, paragraph, list, link etc.), and providing generic page structure (using `<div>`s, now with some ARIA roles sprinkled liberally), but that's its beauty; it's what makes it so universal and accessible.

Digging into the details of HTML5's markup reveals yet another mixed bag. Some interesting inclusions, some baffling ones, a lot of squabbling over some incredibly minor issues, and a lack of a coherent vision to really take markup, and the web, forward.

Then again, criticism of HTML in this regard is hardly new. Here's Clay Shirky in his piece "In Praise of Evolvable Systems" (<http://www.shirky.com/writings/evolve.html>) from—wait for it—1996:

*HTTP and HTML are the Whoopee Cushion and Joy Buzzer of Internet protocols, only comprehensible as elaborate practical jokes. For anyone who has tried to accomplish anything serious on the Web, it's pretty obvious that of the various implementations of a worldwide hypertext protocol, we have the worst one possible.*

*Except, of course, for all the others.*

And it was ever thus.

# THE TRUTH ABOUT HTML5 FORMS

### Put This In Your Form And Smoke It

HTML5 forms are a neat example of HTML5's history. The Architecture Astronauts at the W3C developed an XML-based replacement for HTML 4's forms called XForms (<http://en.wikipedia.org/wiki/XForms>), which was declared a W3C recommendation in October 2004. It was powerful, but completely useless for the web.

In late 2003 an alternative was proposed to extend, rather than replace, HTML 4's forms (<http://hixie.ch/specs/html/forms/xforms-basic-1>). This became known as Web Forms 2.0 from the WHATWG, and was eventually integrated into HTML5.

Note the timeframe: 2003.

The work the WHATWG was doing in the early- to mid-'00s was meant to extend basic HTML forms, and solve the headaches that routinely popped up when using JavaScript to handle complex form interactions.

But in 2005 the idea of JavaScript libraries was gathering steam, with Prototype.js hitting 1.0 ([http://en.wikipedia.org/wiki/Prototype\\_JavaScript\\_Framework](http://en.wikipedia.org/wiki/Prototype_JavaScript_Framework)), and jQuery hitting 1.0 in 2006 (<http://blog.jquery.com/2006/08/26/jquery-10/>).

Web apps were taking off in a big way, and the need for reliable, cross-browser JavaScript libraries was becoming acute. Libraries such as Prototype.js and jQuery (and the many others that have been developed, such as MooTools) met the pressing need, and continued being developed, with jQuery's UI library appearing in late 2007 (<http://blog.jquery.com/2007/09/17/jquery-ui-interactions-and-widgets/>). Other fully-fledged, web app-focused JavaScript frameworks have also emerged (see a comparison here: [http://en.wikipedia.org/wiki/Comparison\\_of\\_JavaScript\\_frameworks](http://en.wikipedia.org/wiki/Comparison_of_JavaScript_frameworks)).

In many ways, these libraries provided all the functionality the WHATWG wanted to build into HTML, and far sooner.

## Going Native Slowly

Native HTML5 form functionality is nevertheless starting to appear in modern browsers, including IE10 (there's no HTML5 forms support in IE9 and below). While the browsers might seem modern, the functionality isn't. By the time IE10 ships it will be almost a decade since Web Forms 2.0 was proposed, and it will be several more years before HTML5 forms functionality becomes mainstream (i.e. when IE10 usage becomes widespread).

That said, more advanced HTML5 form functionality is also starting to appear in mobile. For example, in iOS5 a simple HTML5 element gives us access to the native iOS date picker widget. This highlights the best of web standards—we drop in a simple HTML element, and the browser then provides the appropriate widget. In this case, it's a touch-based widget that would have been unimaginable in the early 00s.

In any case, on the desktop we'll still be using JavaScript for our forms for many years to come, if only to support IE9 and below. The modern JavaScript libraries will become faster, more feature rich, and offer even more functionality (and styling options)—all before the WHATWG's 2003 ideas for forms ever get widespread adoption, such as has been the slow pace of browser development in recent years (Chrome notwithstanding).

Still, there are some handy features we can use today (especially for mobile), and we'll cast a critical eye over the other new functionality to see what can be used now (and whether it *should* be).

## Forms Can Make Or Break A Site

Designers have mixed feelings about forms, ranging from a vague distaste to outright loathing. (Then again, you may be a *form connoisseur* and the exception to the rule.)

Nevertheless, we all need to start loving forms—or at least hating them a little less—because the success of our sites may depend on it. Whether people are trying to sign up, register, check out or even contact us, it could all come down to the quality of the form.

Bad forms are bad business, and oh boy are there some bad forms out there (see: just about every small-time e-commerce site ever). The last thing you want is people willing to give you money falling at the last hurdle because your form wasn't up to scratch.

Designing thoughtful, humane forms (and A/B testing them thoroughly), on the other hand, is good business. Sometimes it's \$300 million of good business: [http://www.ue.com/articles/three\\_hund\\_million\\_button](http://www.ue.com/articles/three_hund_million_button).

I mention this simply because forms appear to be so under-appreciated in the web design world, despite being so crucial to a site's success. People have written entire books dedicated to form design (see: <http://rosenfeldmedia.com/books/webforms/>).

But for now let's focus on HTML5's form features.

## Good News, Bad News

Let's do this good news, bad news style.

The good news is HTML5 has some new form features that will make forms less reliant on JavaScript for common functionality such as client-side validation, range selectors, date widgets and even color pickers.

The bad news is IE9 (and earlier) doesn't support any of them.

The good news is scripts, including jQuery libraries (e.g. <http://flowplayer.org/tools/>) will let us use HTML5 form features where supported, and provide fallback where browsers lack support.

The bad news is some native browser UI's for these form widgets are *worse* than the JavaScript alternatives they'll replace—harder to use, and harder (if not impossible) to style.

The good news is some additions to forms are backwards compatible, and provide some nice touches for iOS and Android devices, so you can use *those* features now.

The bad news is as I write this, large chunks of major HTML5 forms functionality are implemented unevenly, even in non-IE browsers. So we need to tread carefully when implementing them.

Despite all this, we can implement some small things today. So shortly we'll look at HTML5 form features in terms of the 'no-brainers', the kinda/sorta/maybe features, and the features that are interesting but not quite ready for prime time.

## HTML5 Forms Resources

As of writing, tracking down which browser supports which forms feature (and how well they implement it) is a bit of a nightmare. With major browsers on rapid release schedules (Chrome and Firefox, particularly) it doesn't make much sense to document who-supports-what in-depth here, so in this chapter we'll just take a run through the major features, and the major browser support issues. Check out the resources below though if you really want to get your teeth stuck into HTML5 forms.

If you're after authoritative detail on current browser support and browser implementation details for HTML5 form features, see Wufoo's excellent resource: <http://wufoo.com/html5/>. I'll be dropping links to specific pages on Wufoo's HTML5 forms site as we go. They have:

- full demos
- compatibility charts
- screenshots of supported/unsupported behavior
- descriptions of browser quirks
- JavaScript fallbacks, and more.

It's great. Definitely worth checking out.

Other useful resources include:

- The always handy <http://caniuse.com/> has great information on browser support, with links to more information for a given feature.
- Mark Pilgrim's Dive Into HTML5 book has a useful chapter covering some of what's new in HTML5's forms: <http://diveintohtml5.info/forms.html>.
- Peter-Paul Koch has a handy compatibility table of the new HTML5 inputs and form attributes: <http://www.quirksmode.org/html5/inputs.html>.
- There's an extensive (if not particularly reader-friendly) browser compatibility chart on Wikipedia that classes browsers by rendering engine: [http://en.wikipedia.org/wiki/Comparison\\_of\\_layout\\_engines\\_\(HTML5\)#Form\\_elements\\_and\\_attributes](http://en.wikipedia.org/wiki/Comparison_of_layout_engines_(HTML5)#Form_elements_and_attributes).
- Opera's developer site has a brief run down on the new HTML5 form features: <http://dev.opera.com/articles/view/new-form-features-in-html5/>.
- IE10 is currently at its fourth platform preview as I write. The latest release has enhanced HTML5 forms support, so beware of older IE10 information suggesting certain features aren't supported. Check Microsoft's documentation first: [http://msdn.microsoft.com/en-us/library/hh673546.aspx#HTML5\\_Forms](http://msdn.microsoft.com/en-us/library/hh673546.aspx#HTML5_Forms).

## HTML5 Forms: The No-Brainers

HTML5 introduces a few things we *can* start using right away, particularly input field for email addresses, URLs and search terms. These are alternatives to the familiar `<input type="text">`. The good news is browsers that *don't*

recognize these new input types just behave as if the field was just `type="text"`.

HTML5 also introduces a variety of new *attributes* for our input fields (such as `autofocus` and `autocomplete`), some of which we'll look at here, and others we'll touch on below. The attributes here we can generally start using straight away, as browser support is reasonably good, and a *lack* of browser is not particularly consequential.

## New Input Types: Email, URL, Telephone Number, and Search

HTML5 introduces a bunch of new input types, which iOS and Android devices currently use to display a keyboard appropriate for the input type. Sometimes these touches are subtle (the email input type includes the “@” key, the url input type gets a “.com” key, etc.) and sometimes they are more obvious (e.g. the telephone number input `tel` type gets a numeric keypad).

The new input types with useful mobile keyboard variations (for iOS and Android at least) include:

- **Email:** `<input type="email">`
- **URL:** `<input type="url">`
- **Telephone number:** `<input type="tel">`
- **Search:** `<input type="search">`

Email:  URL:



Telephone:  Search:



Figure 8.1. The iOS keyboard variations for the different kinds of input types.

These input types aren't just useful in mobile contexts—they are (search notwithstanding) supposed to provide client-side validation too. So, in Firefox 4+, Chrome, Opera, and IE10, if you use `type="url"` for example, and a user doesn't provide a valid URL, you would get an error bubble something like this (each browser has its own variation):



Figure 8.2. Using the new input types also provides some input validation in supported browsers.

Validation is implemented unevenly for the various input types and across the various browsers (e.g., `tel` has no specified default validation at all), so tread carefully. (Of course, client-side validation is merely a convenience for the user.)

Styling these validation errors is currently highly experimental. There are some experimental CSS3 pseudo-elements in WebKit (for example) that allow you to style the error bubble. You can see the syntax and result at the end of this document: <http://trac.webkit.org/wiki/Styling%20Form%20Controls>.

The search input field is a bit different from the other three we've discussed—the spec doesn't require browsers to do anything special, but some browsers (particularly Safari) round the corners of the search field, and may provide a list of previous searches, and a clear button (a circle with an x in it) when you've entered something.

Older browsers just treat these field as `type="text"`, as mentioned, so there's no harm in using these input types now.

For more, see the Wufoo documentation:

- Email: <http://wufoo.com/html5/types/1-email.html>
- URL: <http://wufoo.com/html5/types/3-url.html>
- Telephone number: <http://wufoo.com/html5/types/2-tel.html>
- Search: <http://wufoo.com/html5/types/5-search.html>

(There are other new HTML5 input types, such as `range`, `number` and `date`, which we'll deal with separately below, in the "I wouldn't yet..." section.)

## Attributes: Autocomplete, Autofocus, Readonly, and Spellcheck

### Autocomplete

```
<input type="text" autocomplete="off">
```

HTML5 specifies an `autocomplete` attribute that is particularly helpful for turning browser autocomplete *off* (autocomplete is on by default). You may want

to do this when a browser's autocomplete suggestions would be inappropriate (e.g. a one-time authorization key, as the spec suggests), or confusing (e.g. autocomplete suggests the user's name when the user should be entering another name). Support has only appeared in recent, modern browsers, and there's no IE support at all. Still, it can't hurt.

## Autofocus

```
<input type="text" autofocus>
```

The boolean `autofocus` attribute automatically assigns focus to a given input when the page loads. The quickest way to see this in action is to go to <http://google.com>—the search box is automatically focused, and you can start typing right away. This is typically done with JavaScript, but it can be a nuisance to some users. For example, with my focus now in Google's search box, I can't use the delete key to go back a page in my history—it thinks I want to delete text in the search field instead.

To deal with this (rather mild) problem, HTML5 moves this autofocus functionality into markup, instead of relying on JavaScript, so your browser can (theoretically) disable it with a preference or extension.

IE9 doesn't support `autofocus`, but Mark Pilgrim details a fallback script here: <http://diveintohtml5.info/forms.html#autofocus>.

## Readonly

```
<input type="text" value="You can't touch this" readonly>
```

HTML5 specifies a widely supported (and self-explanatory) boolean `readonly` attribute.

## Spellcheck

```
<input type="text" name="captcha" spellcheck="false">
```

With the `spellcheck` attribute we can, as per the `autocomplete` attribute, exert some control over default browser behavior. For example, we can turn it

off for inappropriate fields, such as CAPTCHAS. You have to specify whether `spellcheck` should be `true` or `false`.

## HTML5 Forms: The Kinda Maybes

Here are a couple of HTML5 forms features that may help in some contexts, or at least let you experiment on your blog. Browser support here may be mixed; implementations may differ; and fallbacks should be considered.

### Attribute: Placeholder

HTML5 introduces placeholder text for form fields, which is a welcome addition. Designers like it because it lets us put the field labels (or support text) in the field itself and design more compact forms. It's well supported in non-IE browsers (Firefox 3.7+, Safari 4+, Chrome 4+, iOS4+, Opera 11+), and IE10 will support it.

The syntax for this placeholder text is dead simple too. Just add `placeholder="My placeholder text"` to the given field:

```
<input type="search" placeholder="Hit enter to search!">
```



*Figure 8.3. The placeholder attribute in action.*

Neat, right? So why is this in the ‘kinda, maybe’ section?

- **No styling:** Support for styling placeholder text is currently very experimental (see this discussion: <http://stackoverflow.com/questions/2610497/change-an-inputs-html5-placeholder-color-with-css>).
- **No IE9 (and below) support:** The lack of support in IE9 and below is a shame, as this is otherwise well supported. The lack of IE9 support means we'll have to provide alternatives for some time, which can raise a few

tricky issues. Thankfully the Modernizr feature detection script (<http://www.modernizr.com/>) can help in providing a fallback where appropriate.

- **Fallbacks are tricky:** However, fallbacks aren't always appropriate. We could fall back to JavaScript, but JavaScript placeholders for some details (e.g. username and password) can interfere with the browser's built-in autocomplete functionality, which gets ugly.
- **Consistency for conversion rates:** This is the fallbacks-are-better-than-native problem. If (if!) modern JavaScript-based placeholder text improves a form's usability (and therefore its conversion rate, which you can discover through A/B testing if you're really determined) then it should be used for *all* browsers, regardless of their HTML5 support. In fact, if the JavaScript option gives us more design flexibility, why use the native functionality at all? This is true—for the time being—of most HTML5 forms functionality.

HTML5's simple placeholder text may be fine in simple situations (when support is there), but when conversion rates (and design flexibility) are all important, JavaScript solutions are generally going to offer more flexibility (and look a lot nicer to boot).

That's a lot to consider for what is otherwise a simple feature. Hopefully placeholder text support will mature to the point where these issues are moot, but we'll be waiting for IE10 to go mainstream before that happens.

For more, see: <http://wufoo.com/html5/attributes/01-placeholder.html>

## <progress>

```
<progress value="77" max="100">77% complete</progress>
```

HTML5 introduces a `<progress>` element in the forms section of the spec, and it's intended to represent the “completion progress of a task”. It's intended for (surprise) progress bars, and would be commonly (but not exclusively) used in web apps and updated via JavaScript as the task progresses. It could indicate upload progress or, when no value is given, it could indicate that the client is waiting for a response from the server.

We can use the optional attributes *value* and *max* to show progress made so far. We're encouraged to indicate progress inline as text for browsers that don't support `<progress>`.

The idea here is the browser should style this element natively, so it looks like a typical OS progress bar (similar to when you copy a file, for example). As of writing, all modern browsers support `<progress>` excluding Safari, but including IE10 (IE9 and down does not).

Here are a few examples from Chrome on OSX, from Peter Beverloo's demonstration you can try out for yourself (<http://peter.sh/examples/?/html/meter-progress.html>):

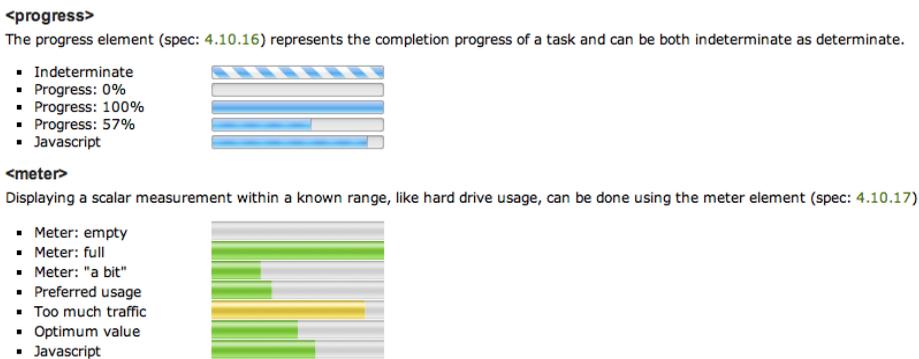


Figure 8.4. Examples of the `<progress>` and `<meter>` elements.

Other browsers just ignore the tags and show the text (e.g. "77% complete"). For the latest on browser compatibility, see Wufoo's handy chart: <http://wufoo.com/html5/elements/2-progress.html>.

This seems like an odd addition on its own, but when you consider the web applications heritage of HTML5 it makes much more sense.

For more, see: <http://wufoo.com/html5/elements/2-progress.html>

### **<meter>**

```
<meter min="0" max="100" value="50">50 of 100 people "liked"
this</meter>
```

While `<progress>` is a ‘kinda, maybe’, `<meter>` should really fall under ‘I wouldn’t yet’, but they make sense together so we’ll keep `<meter>` here.

The `<progress>` and `<meter>` elements sound similar, but they have different use cases and serve different purposes. The `<meter>` element is for gauges, such as a donation gauge indicating \$5,000 progress towards a \$10,000 goal. It could also be used to indicate the percentage of people who voted a certain way, or “liked” something or indicate the number of tickets sold for an event (as the spec suggests), or even represent disk space on a hard drive. It’s explicitly *not* for sole values such as, say, \$5,000 on its own, or for height and weight.

Hickson says he added it to the spec mostly to stop people abusing `<progress>`.

There are six attributes to describe the gauge: `value`, `min`, `low`, `high`, `max`, and `optimum` (only `value` is mandatory, and we’ll touch on some of these attributes below).

Then things get weird.

The `<meter>` element seems simple enough, but things get funky with Chrome (and possibly Safari, if they adopt the experiment WebKit features). Chrome applies native styling, so you get nice meter bars like this:



*Figure 8.5. The `<meter>` element in Chrome.*

However, trying to make things easy and native (drop the element in, the browser does the rest) takes us down a strange path. What if we want to style the `<meter>` element completely differently? Well, we have to undo all the default browser styling, and then apply a lot of experimental CSS3 which has almost no support in modern browsers, as Steve Workman found (see:

<http://www.steveworkman.com/web-design/html-5-web-design/2009/my-problem-with-html-5-styling-meter/>).

How crazy does it get? WebKit includes some experimental CSS pseudoclasses for `<meter>` styling, including `meter::-webkit-meter-even-less-good-value` and even a built-in *star rating system* with `-webkit-appearance: rating-level-indicator`; (see <http://trac.webkit.org/wiki/Styling%20Form%20Controls> for more).

On the one hand, it's nice to see browsers actually doing something with HTML5 elements—a pragmatic reason to use them is welcome. On the other hand, treating `<meter>` like a natively-styled form control takes us down a pretty strange path with lots of bizarre CSS3. Do we really need native star rating systems in WebKit? For now, think of `<meter>` as an experimental novelty, especially as only Chrome and Opera support it.

For more, see: <http://wufoo.com/html5/elements/1-meter.html>

## HTML5 Forms: The 'I Wouldn't Yet But You Can If You Really Want'

Let's look at the `required` and `pattern` attribute, and then several other input types that fallback to `type="text"`, including `number`, `range`, `date`, and `color`.

### Attribute: Required

```
<input type="text" name="musthaveaname" required>
```

The boolean `required` attribute does exactly what you think—it tells the browser that a given `input` (or `textarea`) must have a value before it can be submitted. Note that fields must have a `name` attribute for `required` to take effect.

Safari's half-hearted implementation of this feature (see: <http://css-tricks.com/forums/discussion/11524/modernizr-giving-a-semi-false-positive-with-safari-input->

`attribute/p1`) currently puts this feature in the “I wouldn’t yet...” basket, as it reports to feature-sniffing tools that it *does* support the feature, when it doesn’t really, creating a false positive. This makes it very difficult to create a general fallback without resorting to browser sniffing. IE9 and down does not support the `required` attribute.

Browsers vary in the way they warn a user when they try to submit a form with an empty value for a required field. The ability to style these warnings is very experimental too—it’s the same situation we touched for the validation warnings given by the new input types (e.g. if you enter a non-URL value into a URL input field). WebKit, as mentioned, offers CSS3 pseudo-elements that let us style the error bubbles: <http://trac.webkit.org/wiki/Styling%20Form%20Controls>.

The other big caveat with this attribute is (as the Wufoo page points out) that you only get an error when you submit the entire form. Modern JavaScript techniques check for a value on blur (i.e. as you work through the form), and are therefore more user friendly.

For more, see: <http://wufoo.com/html5/attributes/09-required.html>

### Attribute: Pattern

```
<input pattern="[0-9][A-Z]{3}">
```

The `pattern` attribute allows us to specify a regular expression that a given field’s value must match. (The above regex matches a number followed by three uppercase letters, e.g. 1ABC.) This might be used for ensuring that a user’s post code (or zip code) matches an appropriate format, or a submitted URL matches a particular domain (e.g. it contains facebook.com if providing a facebook profile URL). Regular expressions are not for the faint of heart.

Unfortunately, the `pattern` attribute suffers from the same false positive problem in Safari that we looked at for `required`. It also suffers from the same usability issues we discussed for `required`.

Remember, client-side validation should only ever be used as a convenience for the user, while server-side validation should do the heavy lifting. This sort of

validation is trivial to circumvent, and should obviously never be used for security purposes or input sanitation.

There are also plenty of solid, modern JavaScript validation scripts out there that we're going to be relying on for some time, and offer more user-friendly features (such as validating while you type, or at least as you move through the form).

For more, see: <http://wufoo.com/html5/attributes/10-pattern.html>

### Input Type: Number (Spinner)

```
<input type="number" name="itemquantity" min="2" max="12"
step="2">
```

Earlier we looked at input types for email addresses, URLs, and telephone numbers. HTML5 also introduces an input type for plain old numbers. Desktop browsers generally use this input type to provide a UI for incrementing a field's numerical value (e.g. the quantity of an item in a shopping cart).

As of writing, only IE10, Safari 5+, Chrome, and Opera support this input type. iOS4 just gives you a numerical keyboard, and Opera 11 lets you type any character. The `number` input type accepts attributes `min` and `max` to constrain the range of possible values, and `step` to increment in certain amounts (e.g. by two if you're buying things which only come as a pair).

Browser validation for this field is quite a mixed bag (see the Wufoo page for more). The real problem though is the UI. In WebKit browsers, for example, it's pretty diabolical. They give you this measly "number spinner", and I don't see how older people (at the very least) would cope with such tiny buttons:



*Figure 8.6. The number input type usually gives a number spinner.*

This is a case of the browser makers letting the side down with lousy native interface widgets. There are better JavaScript methods for giving users up/down arrows to increment a value, and any widget you design yourself is going to be more usable than the WebKit implementation.

The UI problems, poor browser support, and inconsistent implementation make it too early to implement.

For more, see: <http://wufoo.com/html5/types/7-number.html>

### Input Type: Range (Slider)

```
<input type="range" name="myslider" min="0" max="10"
step="2">
```

The `type="range"` input gives us a slider, which is fine. Support is absent in Firefox, but has been present in Opera and WebKit (i.e. Chrome and Safari) for a long time, and support will be present in IE10. It's also supported in iOS5. You can also use the attributes `min`, `max`, and `step` to constrain the possible values, and the increments the slider can move in.



*Figure 8.7. Chrome's range slider.*

Browser implementation and styling, however, is a bit all over the place. In my view better jQuery options exist that provide cross-browser support, more features, and a better, more consistent UI.

You could fall back to the native range widget where browsers support it, but why would you bother? It only makes sense to use native widgets if they're actually better (i.e., demonstrably more people actually complete the form)—which is something you shouldn't assume without the conversion rate data to back it up. For now (and the foreseeable future), use JavaScript.

For more, see: <http://wufoo.com/html5/types/8-range.html>

### Input Type: Date (Time/Calendar Widgets)

```
<input type="date">
<input type="month">
<input type="week">
<input type="time">
<input type="datetime">
<input type="datetime-local">
```

HTML5 specifies several date- and time-related inputs (`date`, `month`, `week`, `time`, `datetime`, `datetime-local`) that should bring up either a date picker (for `date`, `month` and `week`) or a number spinner (for `time` values).

Unfortunately, browser support for these is worse than any other input type we've looked at so far. Currently, Opera is the only browser to implement a date picker, with, shall we say, a *functional* look about it.

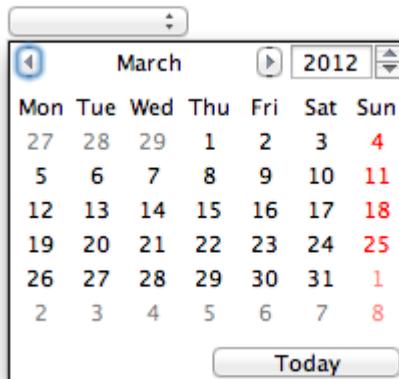


Figure 8.8. Opera's rather functional date widget.

Actually, make that *was* the only browser to support the date input. iOS5 has introduced support for some of these date inputs, giving users the native date picker control to use for these fields, which is really very handy. This kind of

standards support in mobile is the biggest counterpoint to my general “Just use JavaScript” stance in this chapter.

As mentioned, this is web standards at their best—a HTML feature conceived long ago is implemented in a clever way on a platform that didn’t exist when the feature was dreamt up.



Date:



Figure 8.9. The iOS5 date picker.

Should you use the native functionality where possible? For IOS5 users, sure. For Android and other mobile users? When they catch up. On the desktop? In several years time a universally-supported date picker widget will be handy in

non-critical situations. But when the forms are crucial to a site's business, it's a lot of control to give up in terms of styling and custom functionality.

In comparison, jQuery UI (for example) already provides customizable, themeable date picker widgets that offer multiple months, inline display, keyboard shortcuts and more.



Figure 8.10. The jQuery UI date widget is a much more flexible option.

The overwhelming lack of browser support means it's JavaScript or bust for date widgets for the foreseeable future, and there's nothing wrong with that.

For more, see: <http://wufoo.com/html5/types/4-date.html>

### Input Type: Color (Color Picker)

```
<input type="color">
```

For some reason, HTML5 also specifies a color picker.

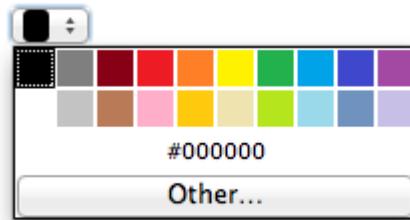


Figure 8.11. Opera's color picker widget.

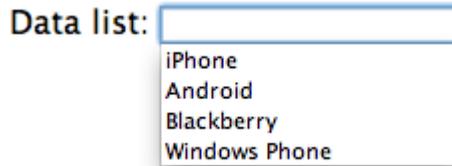
As of writing, only Opera 11+ and (oddly enough) the Blackberry browser support it. WebKit support is currently experimental. There are many better JavaScript alternatives.

For more, see: <http://wufoo.com/html5/types/6-color.html>

### Input Type And Element: Datalist

```
<input list="mydatalist" name="phonenumber">
<datalist id="mydatalist">
  <option value="iPhone">
  <option value="Android">
  <option value="Blackberry">
  <option value="Windows Phone">
</datalist>
```

HTML5 introduces a `<datalist>` element that's used in conjunction with the `list` input attribute to provide a set list of suggestions in a dropdown menu as you type. (As you can see above, the `list` attribute on the `<input>` element matches the `id` on the `<datalist>` element.) These are just suggestions—users can still enter whatever they want. It's essentially just a simple autosuggest feature, and could be useful for, say, choosing your country when you're filling out your details on an e-commerce site. (The usual alternative—a giant `<select>` list—generally proves pretty unwieldy.)



*Figure 8.12. Opera's datalist implementation is pretty good.*

Unfortunately, a complete lack of WebKit support (i.e. Chrome and Safari) makes it hard to recommend at the moment. Firefox 4+, Opera 9+ and IE10 support it, however.

There's also much better JavaScript approaches for this kind of functionality. See, for instance, the very cool Chosen from Harvest: <http://harvesthq.github.com/chosen/>. It offers a variety of `<select>` list replacements, with a very clean UI, solid compatibility for all modern browsers, and graceful degradation for older browsers.

(Wufoo don't mention the poor old datalist, but check the compatibility tables listed at the start of this chapter to see how browser support is fairing.)

## **You Hypocrite. I Thought Requiring JavaScript Was The Worst Thing Ever.**

It may seem strange that I'm advocating JavaScript for forms when I was chastising HTML5 advocates earlier in chapter four for making JavaScript mandatory for basic layout for IE6-8 users. The difference here is that it's still possible to gracefully degrade your form for users with JavaScript disabled—there is no graceful degradation when using HTML5 elements for those users.

## What About Accessibility?

If we're going to use JavaScript for forms, we should still try and ensure our forms are accessible. There's a myth that screen readers can ignore modern, unobtrusive JavaScript and proceed as if JavaScript was disabled. Not true. Here's Roger Johansson ([http://www.456bereastreet.com/archive/201011/accessibility\\_myths\\_in\\_2010/](http://www.456bereastreet.com/archive/201011/accessibility_myths_in_2010/)):

*If screen readers really did not support JavaScript, or screen reader users in general had JavaScript disabled, [then using unobtrusive JavaScript and not thinking much about accessibility would] be a reasonable approach. However, screen readers run on top of web browsers that support JavaScript and, as I mention in "Unobtrusive JavaScript is not necessarily accessible JavaScript", most screen reader users do have JavaScript enabled.*

For screen readers to access forms they need appropriate labeling, descriptions, and structure. (For more information, see this article: [http://webaim.org/techniques/forms/screen\\_reader](http://webaim.org/techniques/forms/screen_reader)) But they still see our JavaScript, so we need to make it accessible to blind users.

And that wraps it up for HTML5 forms!



# THE TRUTH ABOUT HTML5'S CANVAS, GAMING AND FLASH

### Canvas Made Me (Kind Of) Wish I Could Do Flash

Canvas lets us draw programmatically on a specific area of a web page. And it can do some pretty cool things: design enhancements, visualizations, drawing/painting applications, image manipulation, and games (the visual side, at least).

Later we'll get into the nuts and bolts of Canvas. But first let's look at some big picture stuff, and the inevitable comparisons to Flash.

Canvas is not Flash, and comparing a single piece of web technology for (primarily) 2D graphics to a feature rich, broadly supported client environment *and* a mature ecosystem of development tools is a bit apples and oranges. Canvas alongside HTML5 (and related) technologies, however, *is* being touted as a serious Flash competitor. But exploring Canvas (and HTML5 on the whole) made me wish I could do some of the amazing things being done with Flash. This isn't a popular opinion—I'm writing about web standards, shouldn't I hate Flash? It's also not one I thought I'd have—I haven't used Flash for a decade, and as a stereotypical designer/Apple fan (tautology?) I regularly experience the memory-hogging, crash-tastic, resource-consuming, pain in the ass that Flash is on the Mac (and apply Flashblock liberally).

Then I looked at what people were actually building in Flash today, and I felt a twinge of designer jealousy at the sheer *cool* factor of those big-budget, best-of-the-best, runs-anywhere on the desktop Flash sites. Seriously, check out <http://thefwa.com> and sites for games like Lego Star Wars III

(<http://www.lucasarts.com/games/legostarwarsiii/index.jsp>). If there are types of sites Flash is actually good for, it's these. These sites are the benchmark for interactivity on the web. And having studied the current state of HTML5 in terms of interactive web stuff, we've got a long way to go.

## Flash Is Dying, And HTML5 is All We've Got

Nevertheless, it's hard to escape the fact that Flash is dying.

As you know, there's no Flash on iOS devices, as has been widely discussed. (See the April 2010 piece "Thoughts on Flash" by the late Steve Jobs for *why* Apple chose not to support Flash: <http://www.apple.com/hotnews/thoughts-on-flash/>.)

But there won't be any Flash plugin on future Android devices, either. In November 2011, Adobe announced they were discontinuing support for the Flash plugin on mobile devices altogether (see: <http://arstechnica.com/gadgets/news/2011/11/adobe-reportedly-planning-to-gut-mobile-flash-player-strategy.ars>), shifting their focus to Flash-driven native apps and, ultimately, HTML5.

Windows Phone has also never supported the Flash plugin.

The absence of Flash on iOS devices (at a minimum) made offering a mobile, HTML-only version of a site a necessity for most sites. And if you are going to build a HTML version of your site anyway, you need a pretty good reason to build an additional Flash-based desktop site. But the justification could still be made in some cases—a simple HTML site for mobile, and a rich Flash site for desktop, perhaps.

Then Microsoft dropped a bombshell.

In September 2011, Microsoft announced that for Windows 8, the Metro version of IE10 will not support *any* plugins (<http://blogs.msdn.com/b/b8/archive/2011/09/14/metro-style-browsing-and-plug-in-free-html5.aspx>). No Flash, no Silverlight, nothing. Well, nothing in the default Metro interface. Windows 8 will actually have two separate interfaces—a legacy, familiar Desktop mode, and a new, touch-friendly, default Metro mode. If you want to use a Flash site, you'll be kicked out of the

Metro experience and back to Desktop mode, where IE10 will still run Flash and other plugins. That will probably be as clunky as it sounds.

Flash for mobile browsers was a non-starter, but this announcement from Microsoft signaled the beginning of the end of Flash on the desktop, too.

The writing for Flash is on the wall. The default experience of the flagship operating system that will run on hundreds of millions of the world's desktops wont support Flash (or any other plugin technology).

Flash (the plugin) is dying, and I don't say that lightly. Many technology prognosticators have been predicting its demise for some time, but the reality is now undeniable. No iOS support from Apple, no future Android support from Adobe, and now no Metro support from Microsoft.

## Can Canvas And HTML5 Fill The Gap?

Windows 8 will launch in late 2012, and suddenly website owners will be rushing to find HTML5 alternatives to their current uses for Flash. Millions of users will be experiencing a Flash-free desktop experience for the first time.

Consider the implications:

- **Advertising:** Site owners and advertisers aren't going to sit idly by as their ad revenue and click-through rates fall, and they're not going to go back to animated GIFs. There's going to be a huge shift to HTML5-powered (and mobile-viewable) banner ads, and Canvas will have a big role to play there.
- **Media delivery:** Audio and video content will need to be delivered using HTML5, but as we'll see in the next chapter, that's not as straightforward as it sounds.
- **Games:** HTML5 games will also rely heavily on Canvas, as we'll explore later in this chapter.
- **Sites:** Finally, there's going to be a massive number of legacy Flash sites which simply wont work in Metro IE10, let alone on mobile devices. Restaurant sites, I'm looking at you. (There's a niche in converting Flash-based restaurant sites waiting there for someone!) Sadly this also means great interactive sites like the Lego Star Wars III site will not work in Metro IE10, but hopefully this encourages enterprising designers and developers to push web standards to their absolute limits.

## Flash: The HTML5 IDE?

Adobe have been teasing Flash to Canvas export for years—first in 2009, when they demoed such functionality at their MAX conference (see the demo: <http://youtu.be/v69S22ZBBqA>), and again in early 2012 with an announcement of an extension they're working on called “The Adobe Flash Professional Toolkit for CreateJS” (see the announcement and a video of it in action here: <http://blogs.adobe.com/creativelayer/2012/02/28/html5-flash-professional/>). (Erroneous reports did the rounds in mid-2010 that Flash CS5 would export to Canvas, based on the MAX 2009 video re-surfacing, but that was not the case.) The latest demo shows a simple animation making the jump from Flash to Canvas, and playing smoothly on a PC and iPad, but whether this tool will be sophisticated enough for most web professionals (when it's eventually released) remains to be seen.

Flash, the HTML5 IDE, may sound like a contradiction in terms, but so did equating Internet Explorer with cutting edge web standards support just a couple of years ago. Canvas has been around for a long time, and its support in modern browsers is quite mature, so the future of Flash as a HTML5 content creation environment remains solely in Adobe's hands.

(It's also worth noting that In 2010 Adobe also released an experimental FLA-to-HTML tool called Wallaby (<http://labs.adobe.com/technologies/wallaby/>) which relies heavily on SVG and WebKit-only CSS3. And in 2011 Adobe released a preview of Edge (<http://labs.adobe.com/technologies/edge/>), a simple HTML animation tool that relies mostly on JavaScript, despite its “HTML5” claims. However, Adobe says “We've heard loud and clear from the community about their desire to animate content within canvas and SVG” using Edge, and “there are already implementation requests for these features in the Edge product backlog” so Canvas-based animation may arrive from Adobe in a variety of forms. The question is whether Adobe does more than dip their toes in the HTML5 water.)

## And Then Apps Happened

Due to the immaturity of some important HTML5 features (especially audio and video, as we'll see in the next chapter), and the lack of mature design tools for things like Canvas, some site owners may have no choice but to stick with Flash over the next few years, especially as a legacy technology. Or, they'll start pushing apps.

We're entering a strange period for the web.

On the one hand, as far as technology goes, web standards have won, and won handsomely. They have persevered in the face of numerous challenges from proprietary technology (Flash, Silverlight, and many others who have fallen by the wayside), implementor apathy (e.g. Microsoft letting IE stagnate in the 00s), and specification dead-ends (the W3C's XHTML 2 debacle). For web purists, this is an incredible victory that seemed anything but certain several years ago. After all, Microsoft—*Microsoft!*—is releasing a desktop version of its browser that *only* supports web standards, and has not only joined the party in implementing cutting-edge web technology, but is doing it particularly well.

On the other hand, the web standards we've got with HTML5, and the associated development tools, are still not very good. Canvas, as we'll see, has its uses, but Canvas, and HTML5 in the broadest sense, still has a long way to go in both development, adoption, and toolset before it can rival what can be done with Flash today. So what will happen when developers want to do something cool, but they can't do it in the browser?

They'll create apps.

iOS apps. Android apps. Metro apps. Platform-specific apps which, ironically, take us away from the true promise of the web—that it's available to all, from any platform. We saw an explosion in platform-specific software in the 90s as desktop PCs became common, and then web apps started to emerge in the 00s—software that was liberated from any specific platform vendor. Now we have a fierce competitive environment where major platform vendors (i.e. Apple, Google, and Microsoft) see having the best apps as a competitive advantage, *and* two of those vendors are doing their bit to “sunset” Flash as a relevant web technology platform.

This leaves us in the curious situation where there is an exploding market for apps, and a rapidly maturing web standards environment, but some gaping holes in the web technology platform that gets broadly labeled as “HTML5”.

Adobe knows this. I was careful above to highlight the fact they are abandoning the Flash *plugin* for mobile browsers; not Flash on mobile altogether. They are instead focusing their mobile Flash efforts as an environment for native app development (thanks to Adobe’s AIR runtime), and it’s likely we’ll see some of the content that was available on the web thanks to the Flash plugin (especially games and streaming video and audio) move off the web entirely and be reborn as apps, at least until web technology catches up. In the rush to fill the void left by Flash, it’s also possible we’ll see vendor-specific web technology rolled out, and a return to the bad old days of “Best viewed in...” messages. Flash on the web may be dying, but it’s taking some of the web with it.

It’s an interesting time for the web—web standards won, but so far without the technology for the web to go all the way as a ubiquitous platform for all. We’ve instead seen a 1990s-esque explosion of platform-specific apps, and the threat that vendor-specific ‘walled gardens’ will make the web a second class citizen.

All is not lost though—for one, many of those “native” apps are powered in large part by web standards under the surface. And Adobe, for its part, is also investing in HTML5 development tools, and getting involved in the web standards process. The browser makers are improving their browsers at a break-neck speed, and new specifications to bridge the gulf between web standards and native app development are appearing all the time. We’ll explore this more in chapter twelve when we look at HTML5 and web apps.

## Let’s Bury Flash-isms With Flash

If Flash is dying on the web, we should aim to bury some Flash-isms with it. Let’s make sure we remember the lessons of the Flash era, especially when it comes to splash pages, loading screens, pointless animation, annoying widgets, and over-engineered, over-designed “experiences”. They were mostly terrible ideas. Some things simply don’t bear repeating, whether it’s done in Flash,

mountains of JavaScript, advanced CSS3, Canvas, or some ungodly combination of the above.

Let's also be careful of judging a technology too quickly. Canvas is going through its awkward adolescence at the moment—full of potential, embarrassing mistakes, experimentation, monosyllabic grunts, and finally (hopefully!) maturity. Whenever new technology enters the web scene, it's generally shown off in a bunch of experimental or inappropriate ways as a gimmick before settling into its groove and being used moderately and appropriately. Hopefully Canvas finds its groove quickly.

## We're Not In Canvas Anymore

That's the background. Now let's look at what the Canvas element actually does.

The `<canvas>` element defines a bitmap area—or “canvas”, if you will—that you can program and draw on with Canvas' JavaScript API. The Canvas element has been kicking around since 2004, and was rolled into HTML5. All modern browsers support it natively (including IE9, although IE6-8 need to use emulation), as do modern mobile browsers.

The actual element looks like this:

```
<canvas id="mycanvas" width="500" height="200">
  Sorry, your browser doesn't support canvas.
</canvas>
```

Like other things in HTML5, browsers that don't understand `<canvas>` tags just see them as a generic element (like `<mymadeupelement>`) and ignore them, exposing the text between them.

That's it as far as the HTML goes. Everything else happens through the JavaScript API, which lets us:

- draw shapes, gradients and shadows
- manipulate images and text
- create animations (by redrawing the canvas enough times per second).

Working with Canvas is like drawing on a single layer in Photoshop. You only have one layer of pixels to work with, and once you draw over them, they're gone. So to do animation (in games, for instance) we need to redraw the canvas for every frame. Canvas has no sense of managing and manipulating objects (that's more SVG's thing, which we'll look at in chapter eleven), but a variety of libraries (for visualizations and games, in particular) have sprung up to help deal with this.

Given Canvas is manipulated through its JavaScript API, the extent you'll want to get your hands dirty will depend on your interest in JavaScript and drawing graphics programmatically. Here's an example of how we would draw a basic square (using the above `<canvas>` element and `<body onload="draw();">`):

```
function draw() {
    var canvas = document.getElementById('mycanvas');
    var context = canvas.getContext('2d');
    context.fillStyle = "rgb(200,0,0)";
    context.fillRect (10, 10, 100, 100);
}
```

This function gets our `mycanvas` Canvas element (using the 500 x 200px example we looked at above), sets the fill color to red, and then draws a solid red rectangle using the `fillRect(x, y, width, height)` function. As you can see below, we've drawn a red square that is offset 10px on the *x* axis, 10px on the *y* axis, and is 100px wide and 100px high. (I've added a 1px border around our `<canvas>` with CSS so you can see the size of the element itself.)

## Generic Canvas



*Figure 9.1. Our very simple Canvas example.*

We won't delve into the workings of the Canvas API, as there's plenty of solid resources online, including:

- Mozilla's Developer Network Canvas tutorial is a great place to start, as it covers the basics, and has a bunch of links in each section to other resources: [https://developer.mozilla.org/en/Canvas\\_tutorial](https://developer.mozilla.org/en/Canvas_tutorial).
- Opera has a short introduction to the basics of Canvas here: <http://dev.opera.com/articles/view/html-5-canvas-the-basics/>.
- Mark Pilgrim's Dive Into HTML5 has a lengthy chapter on getting started with Canvas: <http://diveintohtml5.info/canvas.html>.
- There's a tutorial for creating a breakout clone with Canvas here: <http://billmill.org/static/canvastutorial/>.
- There's a whole site dedicated to Canvas tutorials here: <http://www.html5canvastutorials.com/>.
- The HTML5 spec for web developers (without all the implementer detail for browser vendors) has a concise run through of Canvas' features here: <http://developers.whatwg.org/the-canvas-element.html>.

Canvas has great potential for gaming and visualizations, along with more mundane uses such as creating charts and even tooltips. But perhaps the most exciting use for Canvas is bringing 3D to the web in a round-about way with WebGL.

We'll look at that in a moment. First, let's look at some examples of Canvas in action.

## Cool Things With Canvas

You can do a bunch of cool things with Canvas, from animations to full blown games. Let's start off with something a bit more modest though: tooltips.

## Tooltips

Can the humble tooltip demonstrate that Canvas can be a better supported option than cutting-edge CSS3?

### Tipped

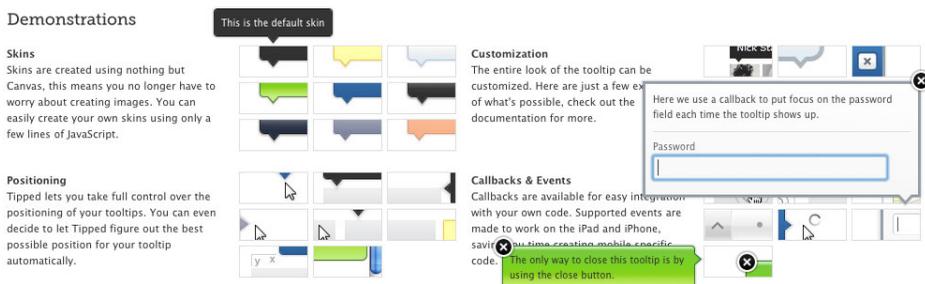


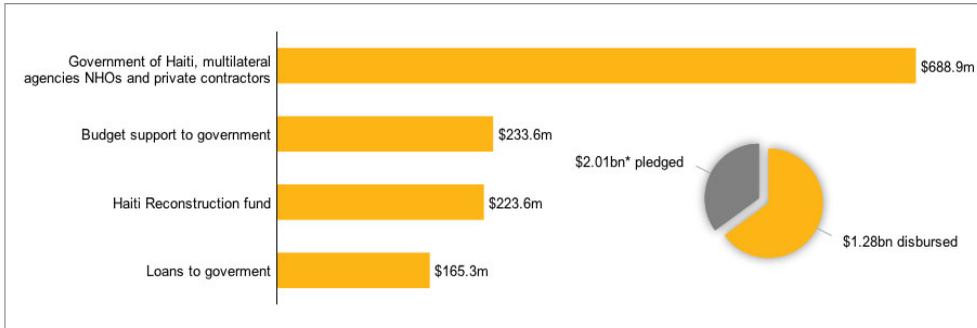
Figure 9.2. Tipped examples.

I think so. Tipped (<http://projects.nickstakenburg.com/tipped>) is a great example of using Canvas to enhance a page. By drawing tooltips programmatically through the JavaScript API, there are no images to worry about. It's easy to create new skins and themes, as well as effects such as rounded corners, shadows, and gradients, all on the fly with Canvas and its JavaScript API. Plus, with IE6-8 compatibility provided through ExCanvas (which we'll look at shortly), we get all the CSS3-style effects with full IE support.

## Charts

Later we'll touch on some SVG-based charting tools (including gRaphaël and the excellent Highcharts). But there's no shortage of Canvas-based charting options either. Here's a small selection.

### RGraph



*Figure 9.3. An RGraph chart.*

Here's a Canvas chart built with the powerful-if-not-all-that-pretty RGraph (<http://www.rgraph.net/>). The beauty of Canvas-based charts is the solid support in iOS (and Android), where Flash isn't an option. (The paid <http://www.zingchart.com/> does Canvas, Flash and SVG if you need better cross-platform support. If you're after something simpler Flot is a popular, free, jQuery and Canvas based charting tool: <http://code.google.com/p/flot/>.)

## Visualize

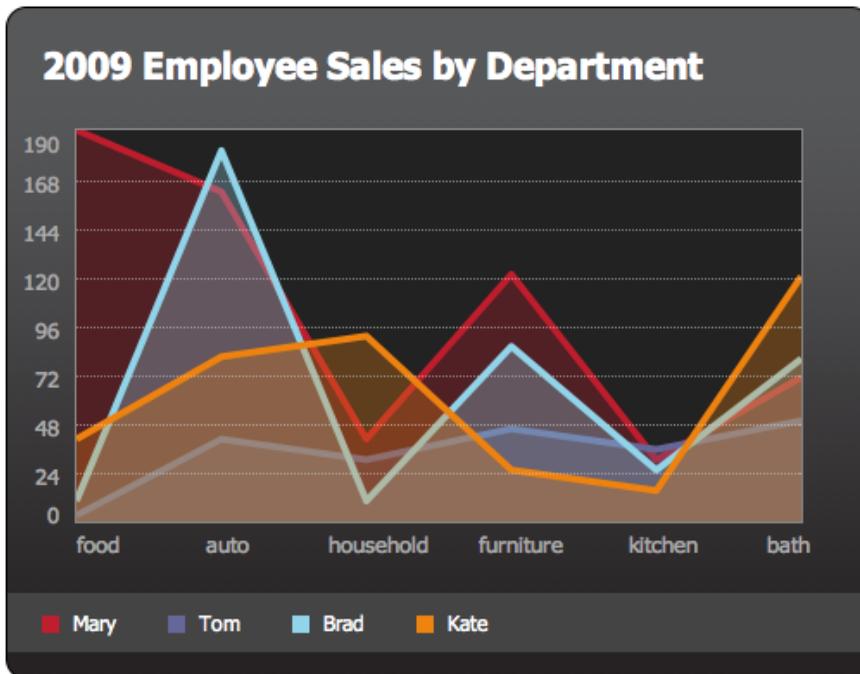


Figure 9.4. Visualize in action.

Filament Group also have an accessible Canvas charting solution with their Visualize plugin, which gives the above result above. It's discussed in detail on their site: [http://www.filamentgroup.com/lab/update\\_to\\_jquery\\_visualize\\_accessible\\_charts\\_with\\_html5\\_from\\_designing\\_with/](http://www.filamentgroup.com/lab/update_to_jquery_visualize_accessible_charts_with_html5_from_designing_with/).

## HumbleFinance

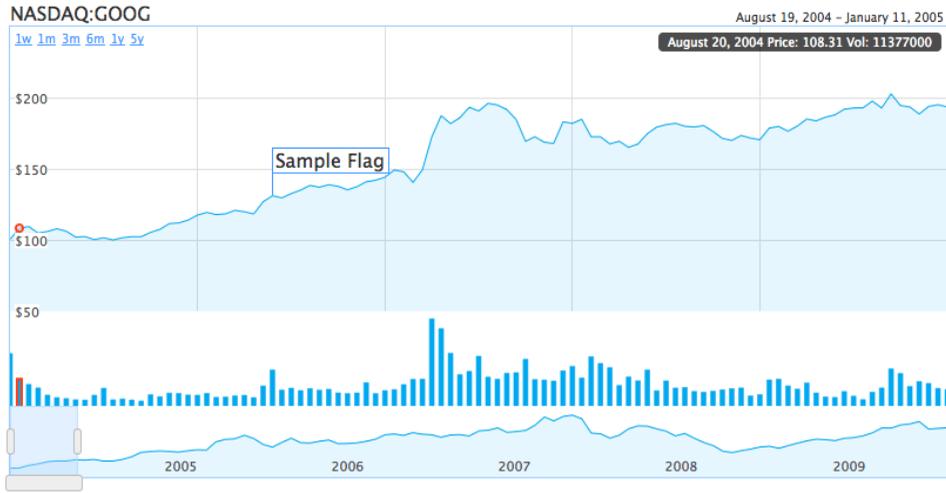


Figure 9.5. HumbleFinance’s sophisticated Canvas-powered chart.

HumbleFinance (<http://www.humblesoftware.com/finance/index>) is a HTML5-driven demonstration of Google Finance-style charting. Because Canvas is just another HTML element, you can easily position other `<div>`s (or any other DOM object) on top of it, which HumbleFinance has done here for the chart labels and other text.

## Peity



### HTML

```
<span data-colour="red" data-diameter="40">1/7</span>
<span data-colour="orange" data-diameter="36">2/7</span>
<span data-colour="yellow" data-diameter="32">3/7</span>
<span data-colour="green" data-diameter="28">4/7</span>
<span data-colour="blue" data-diameter="24">5/7</span>
<span data-colour="indigo" data-diameter="20">6/7</span>
<span data-colour="violet" data-diameter="16">7/7</span>
```

Figure 9.6. Peity examples.

Peity (<http://benpickles.github.com/peity/>) is a jQuery plugin that turns an elements content into a mini pie, bar or line chart. It takes the values in an element like `<span class="line">5,3,9,6,5,9,7,3,5,2</span>` and converts it into a `<canvas>` element that renders the appropriate chart. jQuery Sparklines (<http://omnipotent.net/jquery.sparkline/>) takes a similar Canvas-based approach, and has even more options.

## Visualizations

### Processing.js

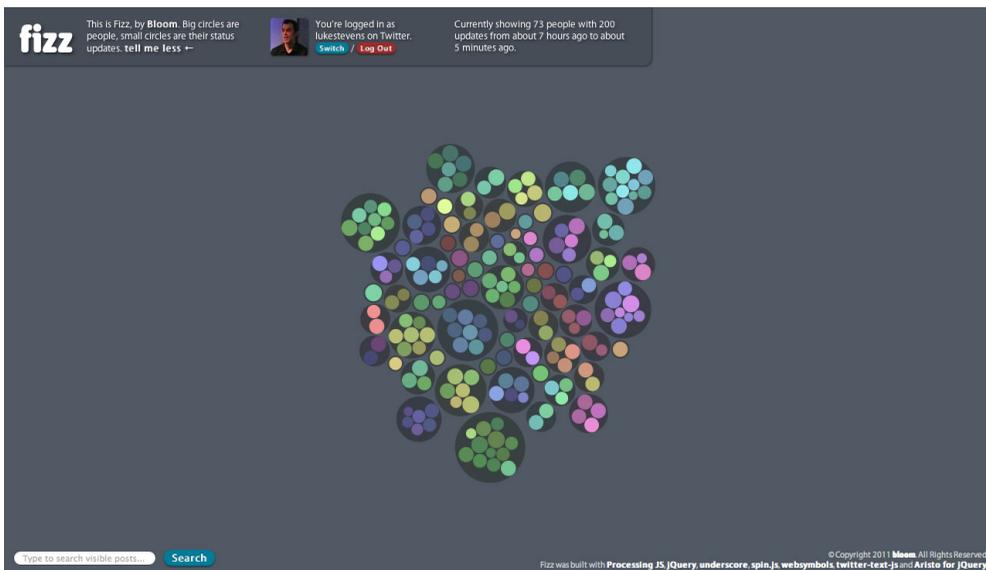


Figure 9.7. Processing.js example Fizz.

Some of the best Canvas examples use Processing.js (<http://processingjs.org/>), the JavaScript port of the Processing visual programming language. Examples range from simple games, to abstract digital art, to visualizations.

## “Evolution of Privacy on Facebook”

# The Evolution of Privacy on Facebook

Changes in default profile settings over time

Trouble seeing the vis? Try switching to [an image-based version](#).

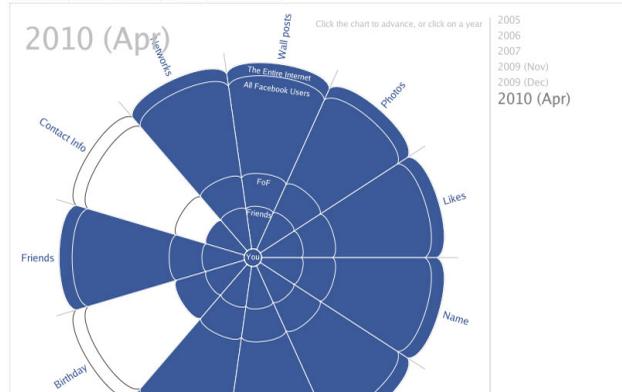
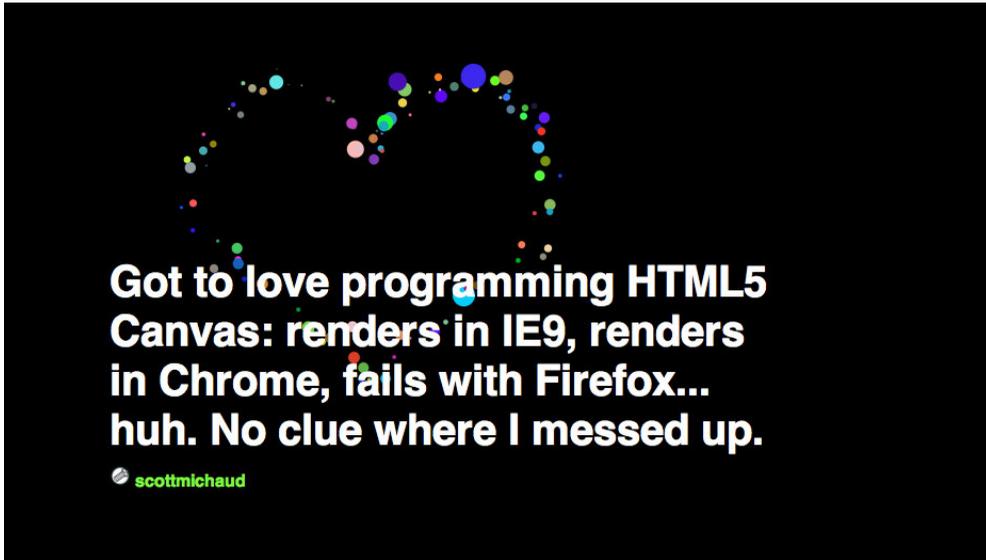


Figure 9.8. The “Evolution of Privacy on Facebook” visualization.

One of the most practical examples of a Canvas-based visualization using Processing.js is the interactive “Evolution of Privacy on Facebook” visualization (<http://mattmckeeon.com/facebook-privacy/>). Because it’s implemented in Canvas it will work on iOS devices, but we still have to worry about compatibility with the (currently) larger IE6-8 group.

## Canvas, Twitter, And Audio Mashup



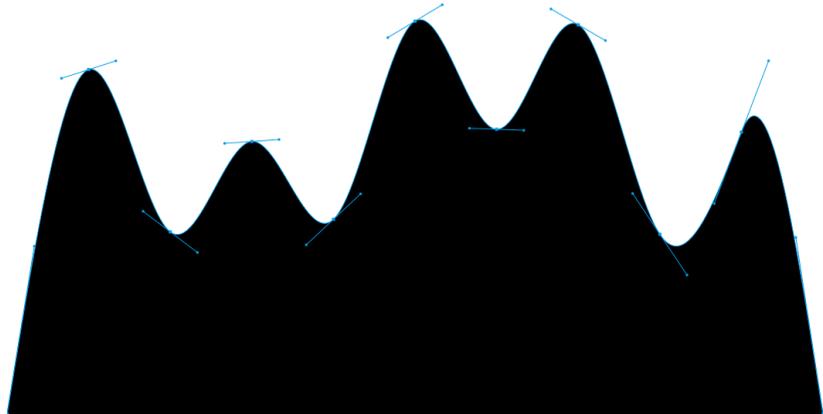
*Figure 9.9. This mashup pulls in HTML5 related tweets, some of which are quite appropriate.*

From the practical to the... well, pretty. This HTML5 Canvas experiment uses Processing.js for particle rendering, and the `<audio>` element to play music (but it's not an audio visualizer). See it for yourself: <http://9elements.com/io/projects/html5/canvas/>. The particles are actually 100 HTML5-related tweets, with their contents rendered as normal HTML in the document. (See the write up here: <http://9elements.com/io/?p=153>.)

## Paper.js

### Paper.js

- About
- Features
- Examples
- Tutorials
- Reference
- Download
- Donation
- License
- Mailing List
- Follow on Twitter
- Watch on Github



*Figure 9.10. Paper.js looks great in motion—be sure to check out the examples on the site.*

Processing.js isn't the only game in town. Paper.js (<http://paperjs.org>) dubs itself as “The Swiss Army Knife of Vector Graphics Scripting” and demonstrates how the bitmap-based Canvas element can be used for advanced vector graphics scripting, complete with a “Document Object Model for vector graphics” and keyboard and mouse interaction. See their examples for more: <http://paperjs.org/examples/>.

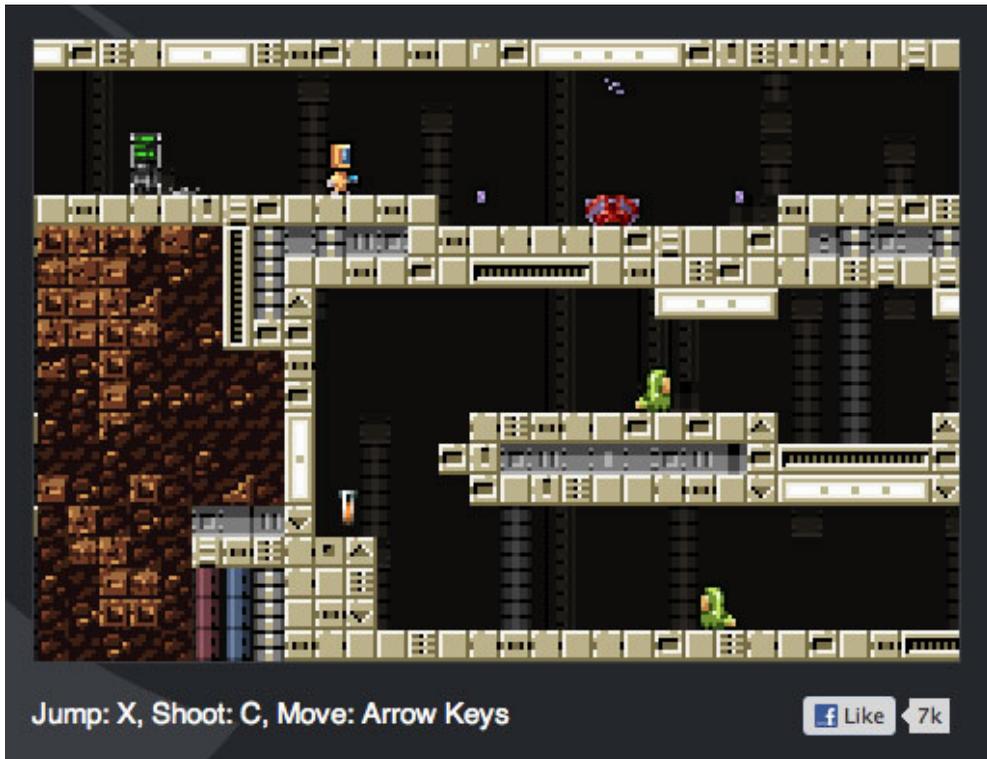
(Smashing Magazine also published a typically extensive comparison of Processing.js, Paper.js, and the SVG-based Raphaël:

<http://coding.smashingmagazine.com/2012/02/22/web-drawing-throwdown-paper-processing-raphael/>.)

## Games

A variety of (mostly retro) games have been built with Canvas. We'll take a look at a handful here, and then look at some amazing WebGL-powered games below.

## Biolab Disaster



*Figure 9.11. Biolab Disaster is a fun little platformer.*

Biolab Disaster (<http://playbiolab.com>) by Dominic Szablewski is a neat example of a retro gaming using `<canvas>` for the visuals. It's a fun little platform game where you run, jump and shoot your way through the levels.

## Canvas Rider

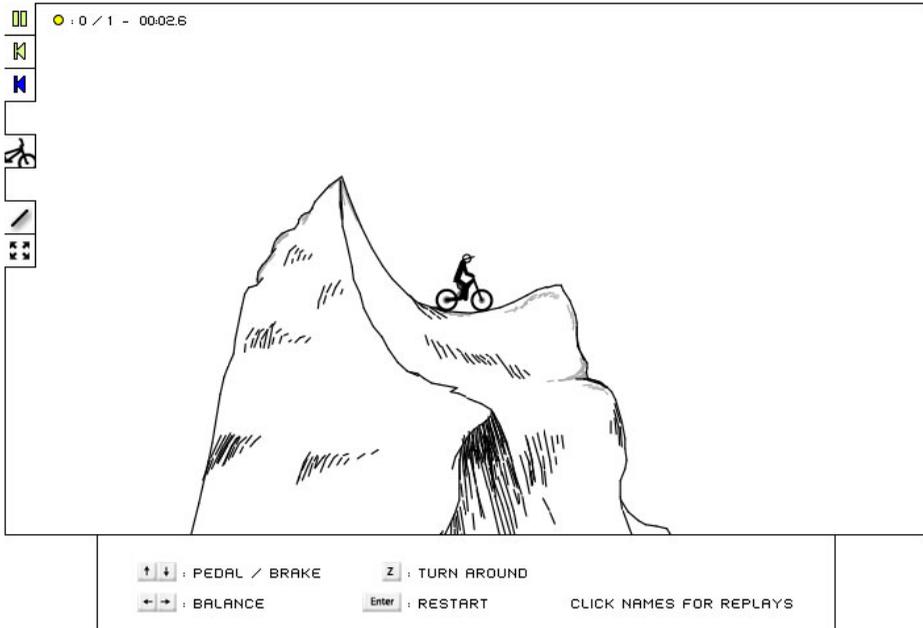


Figure 9.12. Canvas Rider is *addictive*.

Canvas Rider (<http://canvasrider.com/>) is another example of the fun (and hard!) browser games that can be built with Canvas. (Warning: it's highly addictive.)

## Cut The Rope

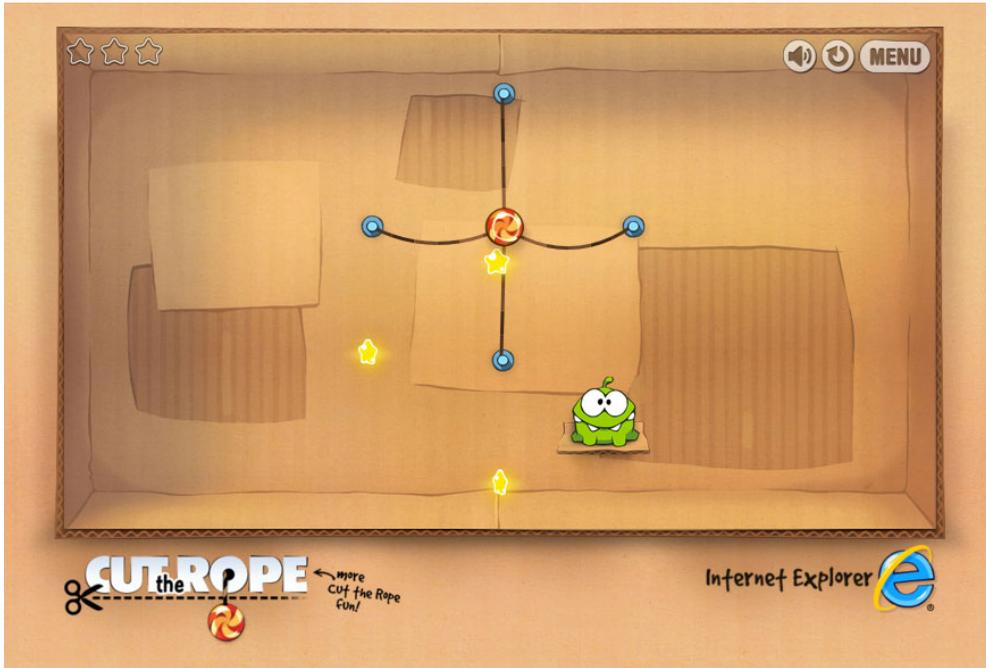


Figure 9.13. *Cut The Rope* takes Canvas-based graphics to a new level.

From small gaming experiments to huge, international hits. The extraordinarily popular mobile game *Cut The Rope* was ported from the original iOS code to HTML5 and released in early 2012. The project was sponsored by Microsoft to show off the HTML5 capabilities of IE9, including its hardware-accelerated Canvas implementation. You can play it in your browser right now:

<http://www.cuttherope.ie/>.

This project demonstrates the potential of web standards for gaming: develop with tools like Canvas, and then easily implement your game in modern browsers, and bundle it up as a native iOS, Android, Windows Phone, and/or Metro app—basically, wherever HTML5 is well supported. We'll take a closer look at gaming and Canvas later in this chapter.

## Imagine Manipulation

### PaintbrushJS



Figure 9.14. PaintbrushJS can perform some impressive Photoshop-like effects.

With Canvas we can do some pretty impressive image manipulation, as Dave Shea's PaintbrushJS library ably demonstrates (<http://mezzoblue.github.com/PaintbrushJS/demo/>). PaintbrushJS lets us apply gaussian blurs, add noise, fade to greyscale (or sepia) and more. And it's all done client-side with Canvas and JavaScript.

## Mozilla Image Editor

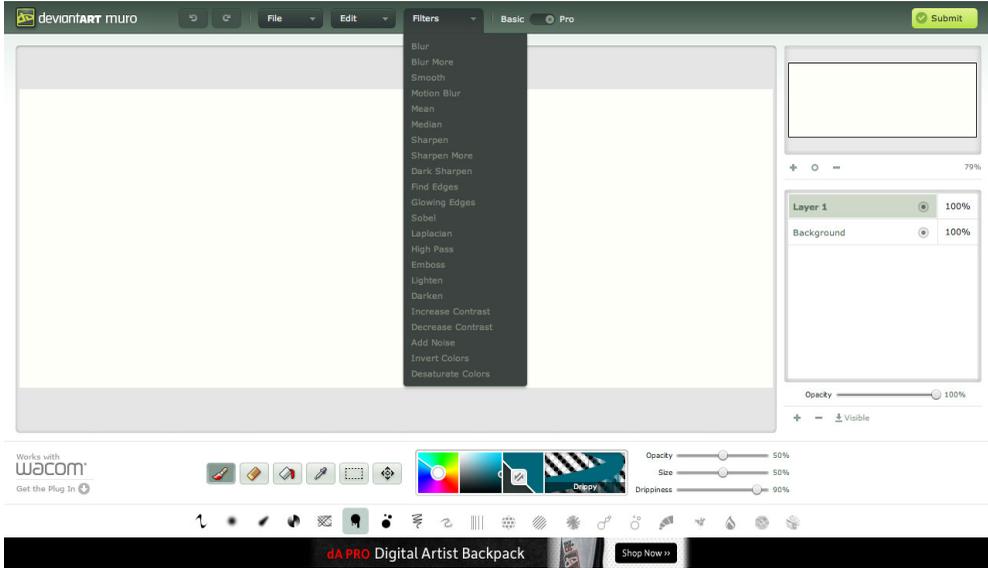


Figure 9.15. Mozilla's image editor and uploader combines a number of HTML5 technologies.

Mozilla pulled a bunch of HTML5 features together to create a HTML5 image editor and uploader (<http://hacks.mozilla.org/2010/02/an-html5-offline-image-editor-and-uploader-application/>). Canvas is used for the image manipulation. I look forward to this sort of functionality being baked right into our Content Management Systems.

## Canvas-Driven Web Apps

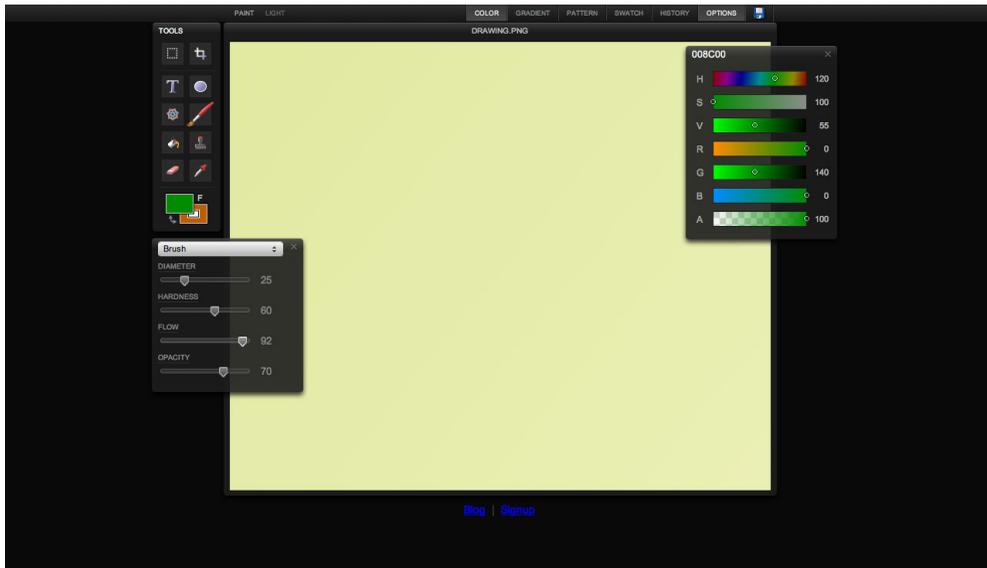
### Muro



*Figure 9.16. Muro is a powerful drawing program, right in your browser.*

Some real, honest-to-goodness web apps use Canvas, which show what's possible in the browser right now. They're mostly painting- or drawing-related (this is Canvas, after all), and none better illustrate this than deviantART's Muro—a free, HTML5-powered drawing/painting app. Try it here: <http://muro.deviantart.com/> or read more about it here: <http://news.deviantart.com/article/125373/>.

## Sketchpad



*Figure 9.17. Sketchpad shows the possibilities of a Canvas-powered painting program.*

Sketchpad is another great HTML5-powered painting app that you can play with here: <http://mugtug.com/sketchpad/>.

## Endless Mural



*Figure 9.18. The Endless Mural site lets you create variations of the above artwork.*

The Endless Mural (<http://www.endlessmural.com/>) is “an interactive, collaborative art website”, powered by Canvas and sponsored by Microsoft for their IE9 launch. The project is the work of Automata Studios and Joshua Davis Studios (Joshua Davis was an early Flash and digital art pioneer). The code has been released as Okapi, an “open-source framework for building digital, generative art in HTML5,” and is available here: <http://okapi.visitmix.com/>.

## LucidChart

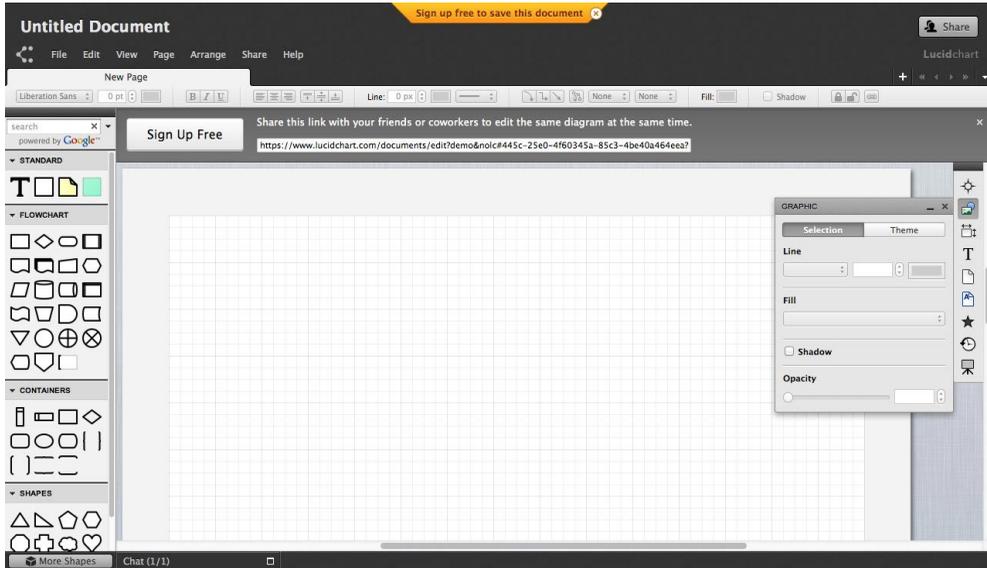


Figure 9.19. LucidChart lets you dive right in with a couple of clicks, so give it a go.

It's not all arty painting though. Some fully-fledged diagramming (paid) web apps also use Canvas, such as LucidChart, which you can test out here:

<http://www.lucidchart.com/>.

## Drawing Interface Elements

### Flash-style Interface Effects

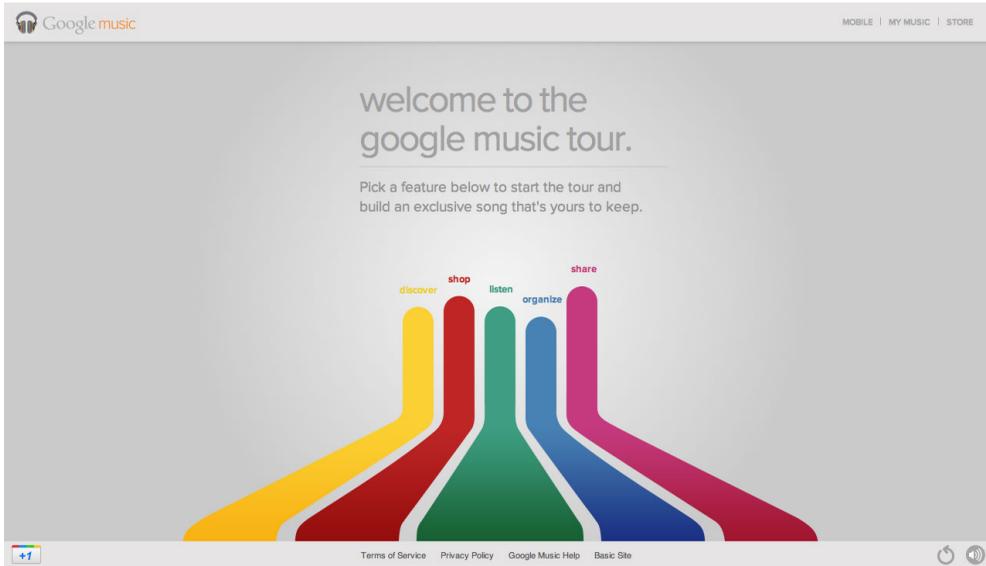


*Figure 9.20. Rally Interactive demonstrate that a clever use of multiple Canvas' can produce impressive effects.*

Rally Interactive pulled off a very impressive animated triangle-to-circle effect to show off their work, with screenshots and statistics pulled from Dribbble. I initially assumed it was a bunch of fancy CSS3 (which would be cool in its own right), but they are in fact using Canvas. See it in action:

<http://beta.rallyinteractive.com/> (and view source to see how they did it).

## Background Animations



*Figure 9.21. The Canvas-drawn background elements on the Google Music site take you on quite a journey as you browse the site.*

Google Music's tour page uses Canvas to render the background animations, which has the thick colored lines scribbling and scrawling as you move from section to section. See it in action: <http://music.google.com/about/tour/>.

## Interface Backgrounds With Liquid Canvas

This demonstration doesn't use any images for the container backgrounds. Each container is laid out via CSS floats and its background is rendered using the operation shown inside. **Liquid Canvas** is based on HTML canvas, jQuery, and excanvas. Tested with FF, IE6(!), Safari, Chrome.

**Liquid Canvas**  
by Steffen Rusitschka

**Example 0**

Plugins are used to perform drawing operations. Here the `fill` plugin gets applied on the `rectangle` plugin. This operation is called application. Gray is the default fill style: a gray rectangle!

```
fill => rect
```

**Example 1**

You can assign properties to plugins, e.g. set a fill color. The syntax is CSS-like.

```
fill(color:#f80) => rect
```

**Example 2**

A plugin exists for rounded rectangles. Oh, btw.: **You should really resize your browser window - now! :-)**

```
fill(color:#fff) => roundedRect
```

**Example 3**

You can also set the properties of plugins on the right hand side. More curves!!!

```
fill => roundedRect(radius:80)
```

**Example 4**

Using an array-like syntax you can chain plugins: First draw a shadow, then fill it. Both plugins in the chain act on a `roundedRect`.

```
[shadow fill] => roundedRect
```

**Example 5**

You can chain as many plugins as you want. Properties can be set on any plugin.

```
[shadow border(color:#fff) fill(color:#f80)] => r
```

**Example 6**

Loads of plugins exist, e.g. the `gradient` plugin - **and you can write your own!** If you are familiar with `HTML canvas` you're ready to go.

```
gradient(from:rgba(255, 255, 255, 0.2) to:#888)
```

**Example 7**

Applications can be chained as well to perform multiple operations inside one canvas. Just separate them with a comma.

```
gradient => rect, border => roundedRect(radius:60)
```

**Example 8**

The application operation can be used recursively by using regular brackets. This gives a different way of chaining operations.

```
fill => rect, mandala => ((border(color:#fff; wid
```

Figure 9.22. The examples may not be pretty, but Liquid Canvas demonstrates some intriguing possibilities.

Here's an idea for us web designers: how about using Canvas to draw the *backgrounds* of interface elements? We could position `<canvas>` elements wherever we like in our markup, even layer them on top of each other as appropriate (using position and the z-index), and then have a small amount of JavaScript render all the graphics—no images required.

This approach could speed up development considerably. No more exporting finicky images from Photoshop to tweak a color scheme for a client. Just change some JavaScript variables and we're done. And with ExplorerCanvas for IE we might even have better browser compatibility than the current state of CSS3. Plus, hardware acceleration is only making Canvas faster on the mobile and the desktop.

The best (and perhaps only) example of this is the Liquid Canvas JavaScript library from 2008. (You can read about it here: <http://www.ruzee.com/content/liquid-canvas>, and see it in action here: <http://www.ruzee.com/files/liquid-canvas/demo.html>.) The demos aren't the prettiest, but the possibilities are certainly intriguing. For example, with Liquid Canvas you can use Canvas to draw

backgrounds behind your divs with shadows, rounded corners, gradients, and stokes. (There's a tutorial available here: <http://www.caffeinedi.com/2009/11/02/using-jquery-and-liquidcanvas-to-add-drop-shadows-borders-rounded-corners-and-other-effects-to-your-website-even-in-ie6-and-ie7/>.)

With a bit more design love, this could be a great way to A/B test certain aesthetic treatments of a layout, all using JavaScript alone—again, no images to export and maintain in CSS. Of course, there will always be limitations with this approach, and presentation should obviously be CSS's domain long term. (And, realistically, most people are going to stick with CSS3.) But Liquid Canvas certainly gives us a very novel approach to drawing interface elements.

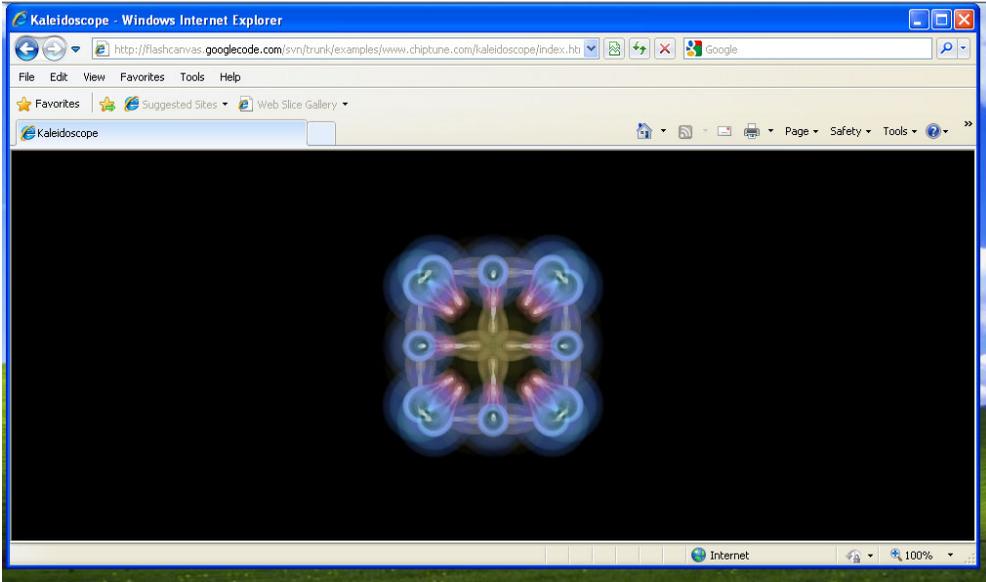
We don't necessarily need to use Liquid Canvas in some (non-IE) situations either. It's been possible to use Canvas elements as CSS backgrounds in WebKit since 2008 (!) as described here: <http://www.webkit.org/blog/176/css-canvas-drawing/>. Firefox 4 added similar support in 2011 (<http://hacks.mozilla.org/2010/08/mozelement/>), and there are workarounds for other browsers (for static Canvas'), as described in the answers here: <http://stackoverflow.com/questions/3397334/use-canvas-as-a-css-background>. With hardware acceleration arriving and the ability to use Canvas as a CSS background, we have some intriguing options for fast, programmatically-generated interface elements.

Consider the possibilities for responsive web design too—we can programmatically redraw CSS backgrounds (in iOS at least) based on the device's resolution. No need to download a desktop-sized image and shrink it down; or maintain different sets of artwork for different devices; just let a script do the heavy lifting for us. (Let me know if you create such a script!)

While we're floating ideas, I'd love to see something like the Photoshop-esque layer style panel for CSS3 effects (which is very cool in its own right—see: <http://layerstyles.org/builder.html>) built for Canvas. It could generate the Liquid Canvas-style JavaScript, which you could then drop into your page and have the effects rendered on just about all browsers. Heck, it would be cool if Photoshop itself could export to Canvas.

Canvas may not be a Flash replacement, but these examples show it certainly opens up some interesting doors.

## The Sometimes Good And Sometimes Bad Canvas Emulation For IE6-8



*Figure 9.23. This kaleidoscope FlashCanvas example demonstrates that impressive Canvas effects are possible even in IE7; they're just slow.*

While IE6-8 don't support Canvas, all is not lost. Several emulation options can give these browsers at least *some* support using IE's native, legacy Vector Markup Language (VML), Microsoft's Silverlight plugin, or Flash.

Each approach has its pros and cons. VML is slow, and loads elements into the DOM. (Canvas elements get recreated as vector elements, so the more that get loaded in the slower it gets.) But animation is smooth, and the Canvas on the whole remains in the DOM.

Flash and Silverlight are both faster. But (as of writing) only around 40% of people have Silverlight installed, and neither Flash nor Silverlight can be manipulated in the DOM.

(There's a more detailed comparison here: <http://uupaa-js-spinoff.googlecode.com/svn/trunk/uuCanvas.js/README.htm>, though it's largely in Japanese, so you'll need to let Google Translate do its thing.)

Throw in issues with interactivity, mixed support for features in the Canvas API, and performance issues for processor-intensive applications, and you have quite a mixed bag. You can still render some impressive animated effects (albeit slowly), or render static images and have them available in the DOM (which is great for charts and the like). But it may be a case of “so close, so far” if performance isn't up to scratch or a crucial feature isn't supported.

Here are the tools that do the emulation. Check out their demos in IE 6, 7 or 8 to get a feel for how they perform (for example, try these in IE: <http://code.google.com/p/flashcanvas/wiki/Examples>). In some cases (such as static Canvas renderings) you may never know it's being emulated, while in other cases you may get away with acceptable-if-not-perfect emulation. Just remember that Canvas emulation can be a pretty murky area—it's rarely a get-perfect-IE6-8-support-for-free card.

Canvas emulation utilities:

- ExplorerCanvas (a.k.a. ExCanvas) is the most well known. It uses VML, and there's an unsupported Silverlight option too. Check it out at: <http://excanvas.sourceforge.net/>.
- FlashCanvas is a Flash-based implementation for Canvas under active development (see <http://code.google.com/p/flashcanvas/> and <http://flashcanvas.net/>.) A Pro version is also available.
- fxCanvas is another Flash-based implementation (albeit less mature) of Canvas for IE, and is also under active development: <http://code.google.com/p/fxcanvas/>
- Finally (and most intriguingly) is the Japanese uuCanvas which claims to offer rendering in either VML, Silverlight, or Flash. Read about it here: <http://uupaa-js-spinoff.googlecode.com/svn/trunk/uuCanvas.js/README.htm>. (Again, Google Translate will help.)

## The Haphazard World Of Web Standards, Or, How Did We End Up With Canvas?

Let's touch on the history of Canvas, as it illustrates how haphazardly these features can be developed. (As was the case with much of HTML5, which is in some ways just a grab-bag of tech that's been around for years.)

*"HTML5" is just a buzzword for 7 years worth of neat stuff.*

—Dave Balmer, <http://www.slideshare.net/dbalmer/rockstar-graphics-with-html5-media-uk>

Do you use OSX's Dashboard feature? That's where Canvas originated, back in 2004. Apple wanted Dashboard widgets to be easy to write, so they based them on the good ol' web stack of HTML, CSS and JavaScript (and native code if you wanted), and used WebKit to render them. (WebKit is the rendering engine behind Safari and Google's Chrome.)

But Apple thought the web stack for rendering Dashboard widgets had limitations, so they added a few features to WebKit, the major one being Canvas. Safari uses WebKit, and so Safari now supported Canvas. And lo, Canvas was on the web was born.

Which, in the tradition of new web technology, caused considerable concern.

Vendors adding browser-specific features to HTML, without any standards process, is what the web standards movement was trying to get *away* from. Had Microsoft and Mozilla responded with their own, incompatible attempts at something like Canvas, we would have a great mess on our hands.

So Ian Hickson *reverse engineered* Apple's Canvas implementation, and put it in the WHATWG Web Applications 1.0 spec. Canvas gained support in Firefox 1.5 in 2005, Opera 9 in 2006, and finally IE9 in 2011. The WHATWG Web Apps 1.0 spec became HTML5 as we saw in chapter one, and Canvas is now an official part of HTML5. Hats off to Hickson for all his work bringing it into the spec.

(For more on Canvas' history, see <http://peter.sh/2010/06/thank-you-microsoft-html5-canvas-is-a-go/>.)

## The Canvas Element And Accessibility

In terms of accessibility, Canvas can be a bit of a nightmare. There's nothing for a screen reader to read—just a black hole and whatever text (if any) the designer has put between the `<canvas>` tags.

Vendors are trying to address this. For example, IE9 exposes the fallback content between the `<canvas>` tags to assistive technology. The idea here is that when the browser *does* support canvas, but the vision-impaired user can't see it, they can still get something useful in the form of alternate text via their screen reader, for example. (Well, that's the theory. At the moment they'll get a lot of erroneous "Your browser doesn't support canvas" messages, because that's how the fallback content has been used to date.) For more see accessibility guru Steve Faulkner's discussion of this feature here: <http://www.paciellogroup.com/blog/2010/09/html5-canvas-accessibility-in-internet-explorer-9/> .

As for the accessibility of the Canvas element itself, it's a thorny issue that has been discussed for *years*, and it still hasn't been resolved. There's a summary of the last few years of discussions here: <http://www.paciellogroup.com/blog/2011/12/html5-canvas-accessibility-discussions-2009-2011/> but that's as far as we've come.

Let's not repeat the accessibility mistakes of decades past with `<canvas>`. We could in theory render fancy text, or an entire web page, or web application (which has been done!), with JavaScript in a `<canvas>` area. But it's a terrible idea, and about as useful (from an accessibility perspective) as having our design as one giant image.

(You can do some crazy things with text in Canvas though, as this tutorial illustrates: <http://www.html5rocks.com/tutorials/canvas/texteffects/>. So we may still do text replacement with such techniques.)

## The Current State Of Canvas

Technologies (web or otherwise) don't exist in a vacuum. Their success often depends on having a conducive environment around them. There's certainly plenty of enthusiasm *for* Canvas, but what about the environment *around* it?

### Primitive Development Environment

It's worth remembering the environment around Canvas is still quite primitive, and so there are few (if any) development tools. (Flash has been successful not just because of the Flash player's ubiquity, but also the mature tools available to developers.)

The situation is changing. HTML5 game engines (e.g. <http://impactjs.com/> which was used for Biolab Disaster) usually have some support for Canvas, but they're pretty niche- and developer-oriented; not quite what we need for general web design.

In terms of tools for designers, there's an Adobe Illustrator to Canvas plugin from Microsoft's Mike Swanson (<http://visitmix.com/labs/ai2canvas/>). And Adobe themselves have been working on a Flash to Canvas export for some time, as we looked at earlier. (They've also been experimenting with more generic "HTML5" export, but they use the term "HTML5" quite loosely. We'll look more at Adobe's position when we discuss SVG in chapter eleven.)

### Performance

One of the exciting things about Canvas is that you can view it on everything from iOS, where Flash isn't an option, to the desktop. In simple, static cases, such as graphs and charts, this is a reality now. But for anything more processor-intensive (such as animation and gaming), recent mobile devices simply haven't been powerful enough to do anything other than the most simple tasks.

This is changing though, particularly as mobile devices get faster, and Canvas gets hardware acceleration (as Microsoft promised and demonstrated for IE9 mobile: [http://www.gsmarena.com/internet\\_explorer\\_9\\_on\\_wp7\\_acces\\_html5\\_drawing\\_test-news-2524.php](http://www.gsmarena.com/internet_explorer_9_on_wp7_acces_html5_drawing_test-news-2524.php)). Apple's iOS5 also dramatically improved Canvas performance, as Grant Skinner documented

in October 2011 (<https://plus.google.com/111971493588974288901/posts/ARX5K8LmNbt> - I've moved the iPad 2 results next to each other below):

*Simple comparison of HTML5 canvas performance on mobile devices, measuring the maximum number of sprite draws per frame at a constant 20 frames per second (ie. 80x80 frames from a sprite sheet image via drawImage):*

*RIM Playbook: 24*

*Desire Z (Android): 60*

*Galaxy Tab 10.1 (Android): 200*

*iPad2 w/ iOS4: 40*

*iPad2 w/ iOS5: 1750*

Sencha also reported dramatic Canvas iOS5 speed improvements:

<http://www.sencha.com/blog/apple-ios-5-html5-developer-scorecard/>. (Sencha's HTML5 blog posts are excellent for keeping up on feature support, by the way.)

Hardware acceleration is making a big difference. But at the moment anything Canvas-intensive can quickly turn into a relative chug-fest on your not-cutting-edge phone or tablet (as the iOS4 and Android device results above demonstrated).

### Limited IE Compatibility

As we saw, IE6-8 can support Canvas to varying degrees and in several different ways (VML, Flash, Silverlight). If you're set on using Canvas this may be a godsend. But compared to Flash it may be a total headache, and limit Canvas' uptake until IE9 becomes the new baseline for design and development. (Remember, IE8 is the end of the line for Windows XP users, so we'll need to wait until Windows XP finally disappears before we can assume native Canvas support.)

As with mobile devices, while simple or static Canvas implementations may be perfectly acceptable with emulation, complex animation and games are probably out of the question. This creates an odd situation when Flash-style interactivity is a necessity—we can either support IE8 and down with a Flash-based experience that won't work on Metro IE10 or mobile, or create something for Metro IE10

and other modern browsers that won't work on IE8 and down. Some poor designers and developers may end up doing both, and may ironically end up creating a more advanced Flash experience for the legacy desktop browsers, and a simpler Canvas-based version for the modern browsers. Here's to the speedy demise of IE8 and the rapid development, adoption, and maturation of new web standards.

## Again With The Glass Metaphor

This is another one of HTML5's glass half-full, glass half-empty situations. It's amazing what people are doing with what was once an OSX Dashboard feature from 2004—from nifty design features (such as tooltips) and interactive experiments to games and full blown web apps. Canvas wasn't designed with these things in mind; it just turned out to be very useful in those situations. That's the glass half full perspective. If you're waiting for a mature, write-once, run anywhere-on-the-desktop environment such as Flash however, it probably looks like the glass is half empty. We're going to be waiting a good while yet.

## HTML5 Gaming: Canvas Or Not?

Canvas often gets mentioned when discussing HTML5 and games, so let's take a brief look at the state of HTML5 gaming. Could you use your existing web skills to write games using HTML and JavaScript that run in any modern browser? Sure, if you're comfortable developing in JavaScript. Will the game be any good? Well, that depends...

One of the biggest trends in recent years has been casual gaming, both in browsers and on mobile devices such as the iPhone. Social games on Facebook have created billion-dollar companies. (Zynga, maker of FarmVille and CityVille, has been valued at up to \$10 billion: <http://www.develop-online.net/news/37066/New-speculation-values-Zynga-at-10-billion>). And casual mobile gaming has been a huge hit too, with the rise of Rovio and its Angry Birds franchise.

So an open, cross-device platform for creating casual, social games is very appealing, from both a hard-nosed business perspective and an idealistic “open

platform” perspective, and HTML5 fits the bill quite nicely. Facebook, for one, is certainly pushing it hard in its developer community.

The key is to understand what type of “gaming” we’re talking about here. Graphically these are often simple 2D games, similar to what was around in the early 90s. In this sense it’s very much “back to the future”—we’re using the latest web tech to create 20-year-old style games in the very modern, mobile world of social networking.

### Is It Even Canvas?

For all the HTML5 hype, some of these HTML games and game engines have *explicitly avoided* features such as Canvas for performance reasons, relying on DOM scripting and CSS3 (which is partly hardware-accelerated on iOS devices) to get the job done. Here’s what one group of developers found from a quick tech demo that followed a HTML gaming engine approach

(<http://sebleedelisle.com/2011/04/html5javascript-platform-game-optimised-for-ipad/>):

*So what’s the answer to getting the performance on iOS? Forget HTML5 canvas, all the moving objects in this game are HTML div elements, we’re just moving them around by controlling CSS properties with JavaScript.*

When discussing “HTML5”, we need to look closely at the technology and techniques people are actually using. What you think is Canvas may well not be. Canvas may be used more widely in web-based gaming as performance improves (and hardware acceleration becomes the norm), but it’s worth keeping in mind just how liberally the term “HTML5” gets thrown around.

Don’t get me wrong: the demos and ports have shown us the scope of what’s possible for HTML5 games on the desktop now, and in mobile browsers in the not-too-distant future. Just keep in mind the phenomena of casual gaming in the browser is less about the latest technology and more about big ideas such as social networking, which the traditional web stack can exploit in pretty interesting ways. That said, hardware-accelerated, 3D gaming is also coming to the browser through Canvas thanks to WebGL, which we’ll look at shortly.

## Getting Started With Canvas Game Development

Nevertheless, if you want to get hands-on with Canvas for gaming check out this tutorial and overview (just ignore the hype in the article):

<http://www.html5rocks.com/tutorials/casestudies/onslaught.html> or this absolute beginners tutorial: <http://www.lostdecadegames.com/how-to-make-a-simple-html5-canvas-game/>. And don't forget the game examples we looked at earlier, including Cut The Rope (<http://www.cuttherope.ie/>) which is probably the most relevant purely HTML5 port to date.

For a detailed run down of all the different technologies available for “HTML5” games (in the broader sense), and their delivery and monetization options, check out this excellent article from January 2012: “The Reality of HTML5 Game Development and making money from it” (<http://www.photonstorm.com/archives/2759/the-reality-of-html5-game-development-and-making-money-from-it>).

## HTML Gaming: Beyond HTML5

There is also a lot of developer interest in taking the web platform beyond HTML5, and including things like the Joystick API, surround sound support, and extensions to CSS. See this detailed W3C write up for more: “Report on the Workshop on HTML.next for Games” (<http://www.w3.org/2011/09/games/>). We'll touch on the post-HTML5 web platform in chapter twelve.

## Canvas: What's In It For Me?

### Canvas For Web Designers

How much Canvas will matter to you depends on where you work and the projects you do. If you work in a big-budget agency, where Facebook components are mandatory for massive, national or world-wide marketing campaigns, you may find new gaming capabilities very interesting.

If you're a freelancer doing client work on tight budgets, the off-the-shelf charting tools like we saw may not be *quite* as sexy, but extremely useful nevertheless. Canvas emulation for IE6-8 may prove very handy as a cross-

device solution (as opposed to Flash-based tools) that covers IE6+ *and* iOS devices.

Canvas-based image editing tools for Content Management Systems could start springing up too, and if you like to tinker, there is enormous room for experimentation. You may want to experiment with rendering interface elements with Canvas (as the Liquid Canvas and Tipped libraries demonstrate), or see how far you can push Canvas in the way studios like Rally Interactive have demonstrated.

### Canvas For Students And Hobbyists

A free, open, and relatively simple platform such as HTML, JavaScript, and CSS could create a fertile environment for kids who want to cut their teeth on simple game design. With tutorials and development libraries springing up all over the place, they have enough information to start making simple (and not so simple) games. It would be great to see this happen in schools, and it wouldn't need much resourcing—just a half-decent PC (or a \$35 Raspberry Pi: <http://www.raspberrypi.org/>) and a modern-ish browser.

### Canvas For Flash Designers

Canvas, along with other not-HTML5-but-cool-anyway features such as jQuery, CSS3 and SVG, may coax more Flash designers to explore HTML and web standards. Flash designers have it pretty good with Flash's mature IDE, advanced features, and widespread support. (I'm sure they'll laugh when slow HTML5 banner ads start appearing too.) Just remember with Microsoft's moves, the writing for Flash is on the wall.

In any case, the more that can be *delivered* in HTML (even if the animations are created in Flash), the better. We need the creativity of Flash designers to push the boundaries of HTML5 and new web technologies as far as possible. Hopefully Adobe does more to deliver on their export-to-Canvas promises and other "HTML5" authoring tools too.

## Suck It And See

We have enormous scope for weaving the `<canvas>` element into our web pages in subtle (or not-so-subtle) ways. But whether Canvas becomes a staple web design tool, or simply the Java applet of our time, is up to us. Let's give it a go and see what we can come up with.

## 2D Canvas' 3D Future: WebGL

I've left the best to last—one of the most exciting developments related to Canvas is WebGL (Web-based Graphics Library). Despite Canvas' ostensibly 2D origins, the new WebGL standard gives Canvas a hardware-accelerated 3D context powered by OpenGL—if the browser (and underlying hardware) supports it. This opens the door to modern, 3D gaming in your browser.

(The WebGL spec wasn't developed by the W3C or the WHATWG, but instead by non-profit technology consortium Khronos Group, which sprang from Mozilla. So it's not HTML5 per se, but it's still cool.)

The WebGL working group includes Apple, Google, Mozilla, and Opera. (See the Wikipedia entry for more: <http://en.wikipedia.org/wiki/WebGL>.) Notice who's missing from the list? Yep, Microsoft. While other major browser vendors are pushing ahead with actual or experimental support, security concerns have given Microsoft *very* cold feet. For example, a post in June 2011 from Microsoft's security research blog titled “WebGL Considered Harmful” said:

*We believe that WebGL will likely become an ongoing source of hard-to-fix vulnerabilities. In its current form, WebGL is not a technology Microsoft can endorse from a security perspective.*

Translation: don't hold your breath for WebGL in IE anytime soon.

This isn't just Microsoft being overly cautious or hostile to competing technology either. John Carmack (id Software founder, and highly-respected gaming developer) tweeted soon after the blog post ([http://twitter.com/#!/ID\\_AA\\_Carmack/status/81732190949486592](http://twitter.com/#!/ID_AA_Carmack/status/81732190949486592)):

*I agree with Microsoft's assessment that WebGL is a severe security risk. The gfx driver culture is not the culture of security.*

Nor is it solely Microsoft who has security concerns. Apple has yet to enable WebGL in the browser on iOS devices for similar reasons (if these comments on Hacker News are accurate: <https://news.ycombinator.com/item?id=3252777>). Apple is, however, enabling WebGL for iAds in IOS 5, so WebGL-powered marketing may be a reality sooner rather than later (for better or worse).

### 3D On The Web: WebGL Alternatives

WebGL isn't the only game in town, either.

Flash 11 launched in October 2011 and recently brought hardware-accelerated 3D to the browser with its "Stage 3D" technology (previously known as "Molehill"). You can see demos here: <http://www.adobe.com/devnet/flashplayer/stage3d.html>. Epic has ported their Unreal Engine 3 to Flash 11, so the gaming potential is certainly there: <http://www.anandtech.com/show/4933/flash-11-supports-unreal-engine-3>.

Microsoft's Silverlight 5, released in December 2011, also introduced hardware accelerated 3D, but given Microsoft's plugin-free plans for Metro IE 10, it seems unlikely Silverlight has a future in the browser, and will probably morph to become a development environment for Metro apps.

Interestingly, Unity Technologies, who make the popular cross-platform (and very mobile-friendly) 3D engine Unity have been distributing their own browser plugin, which they claim has seen 60 million downloads. They're also rolling out an export-to-Flash option so users don't need to download the Unity browser plugin.

So, one way or another we'll be seeing a lot more 3D content in the next couple of years, particularly for games, which still may be compelling enough to get Windows 8 users who use the Metro interface to switch to desktop mode to run Flash. We're in the midst of the 2D browser-based gaming explosion (CityVille hit 100 million active monthly players in 2011), and the 3D browser-based gaming boom is about to follow.

WebGL also has old-school predecessors. Displaying 3D on the web isn't new—VRML (“Virtual Reality Modeling Language”) for example dates back to 1994 (see: <http://en.wikipedia.org/wiki/VRML>). But with hardware-accelerated 3D now pretty much a given on any platform (including smartphones), the potential for 3D on the web for gaming and beyond (e.g. 3D modeled product previews, medical models, maps, and more) is infinitely greater.

Who knows? Maybe we'll finally have the technology for that 3D “virtual” shopping mall experience we've all been, uh, craving...

## Show Me The Demos!

When it comes to WebGL, seeing is believing. So fire up the latest version of Firefox, Chrome, or Opera and check out these cool demos. And remember: this is all happening in `<canvas>`—just another element in the DOM (which we can push around with CSS3).

## Angry Birds

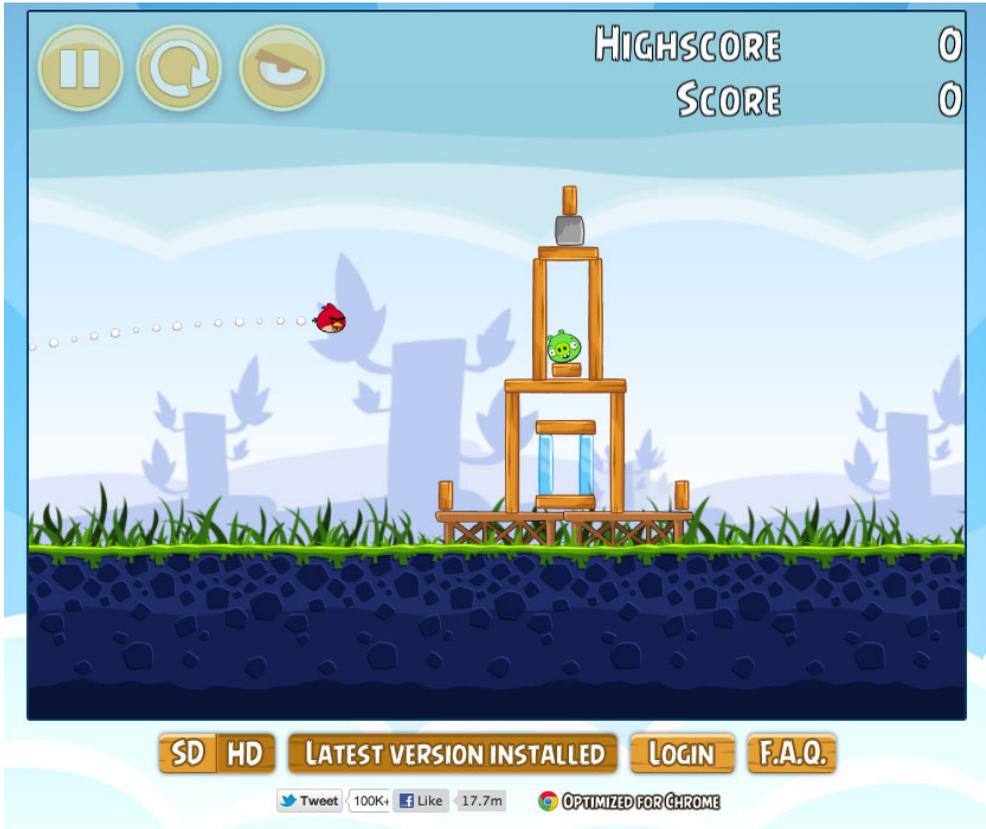
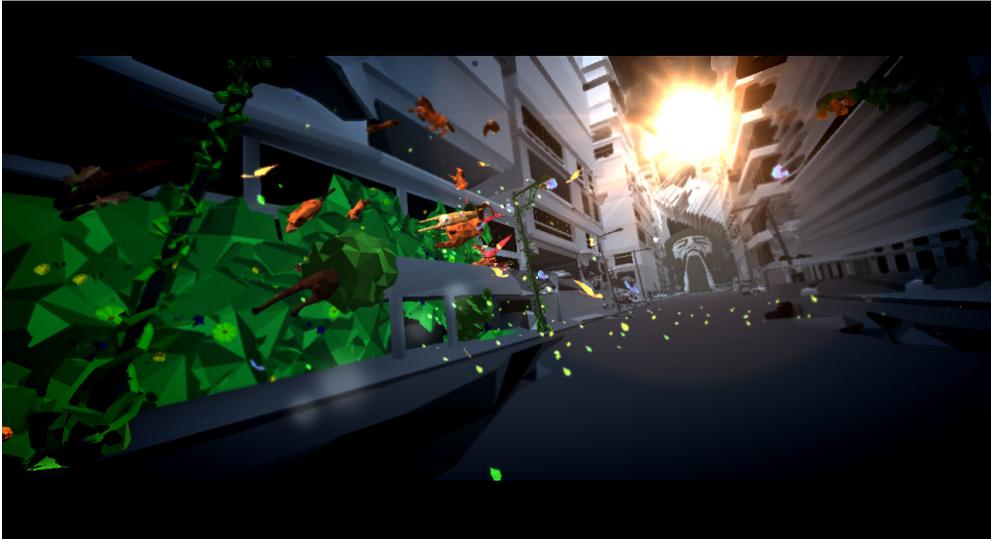


Figure 9.24. The Angry Birds juggernaut hits the web thanks to WebGL.

Yep, Angry Birds in the browser (<http://chrome.angrybirds.com>). WebGL delivers hardware-accelerated 2D graphics as well. (Take a moment to consider the implications of 2D, WebGL powered interfaces for web sites, and not just games.) Try Angry Birds on the web for yourself at: <http://chrome.angrybirds.com>. Interestingly, when WebGL isn't available this falls back to DOM animation (including moving 2D Canvas elements around), so you can compare the performance and how much a 2D game like Angry Birds benefits from hardware acceleration. (You still need Flash for sound, due to the poor state of audio for games in browsers.)

## Rome “3 Dreams in Black” Interactive Music Video



*Figure 9.25. The Rome experience is an absolute must see.*

This incredible music video for Danger Mouse’s Rome project is a great example of how interactive WebGL can be. Check it out here in Chrome: <http://www.ro.me>, though it should be viewable in the latest Firefox versions as well. It’s an amazing experience, and there are even user-submitted 3D models in the desert at the end.

You can watch a video of the team behind the film (though “film” doesn’t do it justice) with clips of the experience here: <http://www.youtube.com/watch?v=ReH7zzj5GPc>.

It will be exciting to see what designers, artists and engineers can produce with this technology in the future. (I just hope it doesn’t turn into too much of a marketing tool. Imagine reading the news on a mainstream news site, when suddenly a big Canvas element is overlaid and you’re thrust into a 3D “brand experience”.)

## glfx.js Image Manipulation



*Figure 9.26. The tilt shift effect in glfx.js is one of many Photoshop-esque effects.*

Earlier we saw how Canvas can manipulate images in 2D, but WebGL unleashes even more power thanks to its hardware acceleration: <http://evanw.github.com/glfx.js/>. The WebGL-powered glfx.js image effects library lets us apply hardware-accelerated, Photoshop-like filters such as brightness/contrast, curves, denoise, hue/saturation, unsharp mask, lens blur, tilt shift, triangle blur, zoom blur, color halftone, perspective transformations, swirls and more. Check out the demo to see it in action: <http://evanw.github.com/glfx.js/demo/>

## Quake II

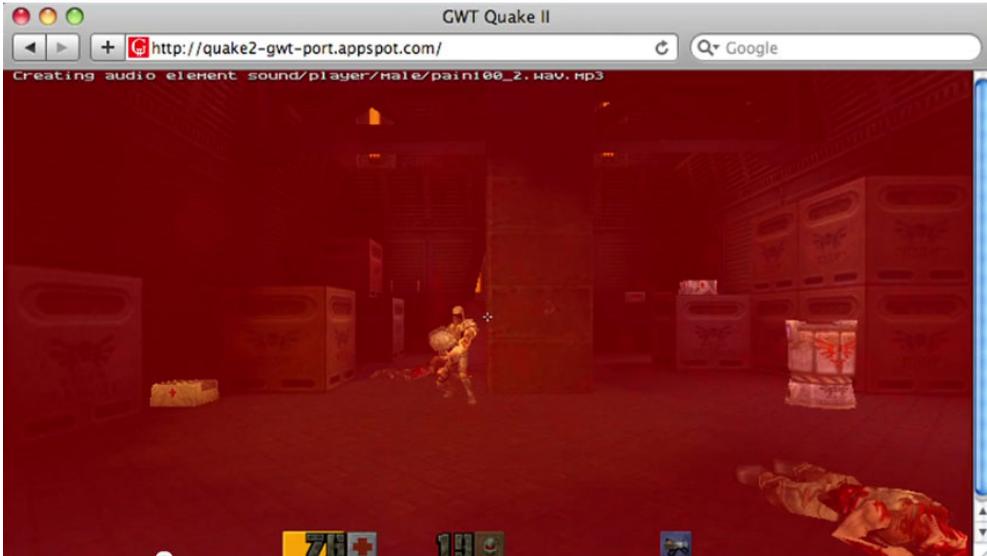


Figure 9.27. Quake II running in the browser using only web technologies.

Quake II has also been ported to WebGL, using “WebGL, the Canvas API, HTML 5 `<audio>` elements, the local storage API, and WebSockets to demonstrate the possibilities of pure web applications in modern browsers such as Safari and Chrome”. For more, see: <http://code.google.com/p/quake2-gwt-port/>. (You need to download and compile the code to actually play it, but there’s a video of it in action here: <http://www.youtube.com/watch?v=fyfu4OwjUEI>. There’s also a Quake 3 demo level in WebGL here: <http://media.tojicode.com/q3bsp/>.)

In a response to a tweet that said:

*Not sure if the best endorsement of JS engine speed in 2010 is ports of games from 1997...*

Joel Webber, one of the engineers behind the port, wrote (<http://blog.j15r.com/2010/04/quake-ii-in-html5-what-does-this-really.html>):

*What's the point? What this code currently proves is that it's feasible to build a "real" game on WebGL, including complex game logic, collision detection, and so forth. It also did an excellent job exposing*

*weak spots in the current specs and implementations, information we can and will use to go improve them. [...]*

*[O]ne can envision a world where game developers can deploy their games as easily as we currently deploy web pages, and users can play them just as easily. And that's a pretty damned compelling world to imagine.*

Send a link. Play the game. That's what WebGL (and the other emerging 3D technologies) will enable. Or, in fact, *is* enabling.

### GT Racing: Motor Academy



*Figure 9.28. Playing GT Racing live in the browser on Google+.*

In December 2011, Gameloft debuted a WebGL-powered version of their game GT Racing: Motor Academy on Google+ (<https://plus.google.com/u/0/games/777131296458>), which runs in Chrome and Firefox. This is an interesting look at not just the future of gaming technology in the browser, but distribution through social networks too. (Gameloft's Baudouin Corman discussed these issues with

Gamasutra here: [http://www.gamasutra.com/view/news/39273/Gameloft\\_Embraces\\_HTML5\\_With\\_3D\\_Game\\_GT\\_Racing.php](http://www.gamasutra.com/view/news/39273/Gameloft_Embraces_HTML5_With_3D_Game_GT_Racing.php).)

## Skid Racer



*Figure 9.29. The console-like Skid Racer may represent one of the first steps in modern, 3D, web-distributed games.*

GT Racing isn't the only WebGL-powered racing game out there either. Skid Racer is an original, WebGL-powered “kart racer” that's available in the Chrome Web Store (<https://chrome.google.com/webstore/detail/bhaojooagiaaiidlnfhkkafjpbnnno>). (The irony of a web-based, Chrome-only game should be noted, but we'll give the developer the benefit of the doubt and put it down to distribution issues.)

## More WebGL Demos

Here are some more WebGL demos and examples:

- Mozilla's **Flight of the Navigator** demo (<http://videos-cdn.mozilla.net/serv/mozhacks/flight-of-the-navigator/>) is an interesting mashup of new web technology. It's mostly a WebGL fly through with HTML5 audio, but with real-time Flickr and Twitter integration. Some of the effects

(interlacing on billboards, ticker-tape displays, audio visualizations) are done with Processing.js and 2D Canvas (which we looked at earlier), and then used to texture the 3D WebGL objects. Crazy. (Read more about it here: <http://vocamus.net/dave/?p=1233>.)

- **Google MapsGL** is a WebGL-powered version of their ubiquitous maps service, where WebGL is used to enhance the experience. Read more about it and watch a demo here: <http://support.google.com/maps/bin/answer.py?hl=en&answer=1630790>.
- **CycleBlob** is a fun lightcycle game on 3D surfaces: <http://cycleblob.com/>
- **Tankworld** is a great little tank shooter: <http://www.playtankworld.com>
- Süperfad's **Mission Control**, their “global traffic visualizer”, is a beautiful 3D visualization of their Google Analytics traffic. See it live here: <http://superfad.com/missioncontrol/> and read about it here: [http://superfad.com/blog/post/mission\\_control](http://superfad.com/blog/post/mission_control).
- For more WebGL demos see these experiments: <http://www.chromeexperiments.com/webgl> or this list of traditional Canvas as well as WebGL games: <http://www.netmagazine.com/features/top-20-html5-games>.

## Still Early Days For WebGL

The WebGL games and demos show incredible promise, but it's not all beer and skittles for WebGL development at present. One developer on Hacker News commented on the hardware incompatibilities, software bugs and inconsistencies that can make it a difficult development environment for small teams, saying (<https://news.ycombinator.com/item?id=3253016>):

*[W]e found so many inconsistencies across different hardware and different browsers that it made it not worth it to work on a WebGL project for the time being, especially for a small team. We wrote a number of runtime checks, but we still could not account for all the bugs, or find ways around every one of them.*

Nevertheless, as the demos above show, amazing things are possible. And if you're interested in the raw performance details of WebGL, check out this extensive post: “HTML5 2D gaming performance analysis”

(<http://www.scirra.com/blog/58/html5-2d-gaming-performance-analysis>) from Scirra, developers of HTML5 game making tool Construct 2, which concludes:

*Hardware accelerated 2D canvases are fast, but WebGL is still way faster. It is in the realm of native engines, which is frankly amazing considering how inefficiently designed Javascript is, and is a testament to the amazing job browser makers have done in writing JIT compilers for Javascript.*

Hopefully Microsoft and Apple find a way to enable WebGL on their desktop and mobile platforms soon, as the possibilities of widespread, hardware-accelerated 3D in the browser courtesy of WebGL are very exciting indeed.



# THE TRUTH ABOUT AUDIO & VIDEO IN HTML5

### Hear No `<audio>`, See No `<video>`

Both `<audio>` and `<video>` are welcome additions to the HTML spec. We don't use proprietary tools (such as Flash) to display images, so why should we need them to play audio or video?

Don't get me wrong: I'm not here to bash Flash. Without it we wouldn't have YouTube, Vimeo and the video revolution that has taken place over the past few years. And Flash still provides advanced video features (such as live streaming, full-screen playback, and DRM, as icky as DRM is) that have either no, or very preliminary implementations as open standards.

Nevertheless, HTML5 `<video>` and `<audio>` have become almost mandatory for media delivery for one reason: iOS. Given Apple's decision not to allow Flash on their mobile devices, the only way to embed video and audio so it's available for iPhone, iPad and iPod Touch users is to use these new HTML5 elements.

But it doesn't stop at iOS. As we saw in the previous chapter, in late 2011 Adobe announced they were abandoning the Flash plugin on mobile altogether, and shifting their focus to native apps and HTML5. Plus, Internet Explorer in Microsoft's new default Metro interface for Windows 8 won't support *any* plugins, period. (For more on the demise of the Flash plugin, check out the discussion in the previous chapter.)

We're racing towards a post-Flash future. Unfortunately, our open technologies are just not ready to replace all Flash offers. Web standards won, and we're going to get caught with our pants down if we're not careful. That leaves apps to fill the gaps, taking us back to the 90s with platform-specific software.

We'll return to the issue of our post-Flash future shortly. For now, let's look at the new HTML5 audio and video elements as they currently stand. While these elements are specified quite simply, issues outside the spec make their implementation... *interesting*, to say the least. First, let's look at the basics.

## Native <video> And <audio> In Action

So how do we use these new elements? In theory, it's pretty simple. Let's start with <audio>.

### The <audio> Element

For audio, we can use:

```
<audio controls autoplay loop muted preload="auto">
  <source src="myaudio.ogg" type="audio/ogg">
  <source src="myaudio.mp3" type="audio/mpeg">
  <!-- Fallback content, such as a Flash player, here -->
</audio>
```

This will give us something like this Chrome example (note that browsers render the <audio> element however they want, and there can be considerable difference in player size):



Figure 10.1. The default Chrome audio player.

Okay, let's run through this. Firstly, you *can* stick the `src` attribute in the opening `<audio>` tag if you like, instead of using the `<source>` element. But due to mind-bogglingly annoying issues with codec support (which we'll delve into shortly) we often need to specify two source files for maximum HTML5 compatibility.

And that's what the `<source>` element is for. Browsers work through the list of `<source>` elements until they either find a file format they support or (if we've included one) they get a fallback option—perhaps a link to the file, or (more commonly) a Flash media player. (See the resources below for a tutorial on implementing a Flash fallback.)

Both `<audio>` and `<video>` are implemented in a backwards compatible way, insofar as older browsers (such as IE8) will ignore the `<audio>` and `<source>` elements altogether (IE8 just sees those elements as generic tags, like `<mymadeuptag>`). This means whatever fallback content we include will be visible to the older browsers, but ignored by the modern ones.

We can also use scripting to give a Flash fallback to browsers that support the `<audio>` (or `<video>`) elements but *not* the codec we've used. We'll look at media players that do the heavy lifting for us at the end of this chapter.

### **<audio> Attributes**

Back to the `<audio>` element. The attributes `controls`, `autoplay`, `loop` and `muted` are *boolean* attributes—including them makes them `true`, and excluding them makes them `false`. (But given that `autoplay` and `loop` are tools of the devil, we should probably always leave them out.) The `muted` attribute makes the browser's player default to mute (though support for this may be patchy), and the `controls` attribute tells the browser to use its native controls. (The elements can also be controlled through the JavaScript API.)

There's also a `preload` attribute that's *not* boolean, but instead can be `none`, `metadata` (preload the metadata for the file only) or `auto`, which usually means the browser will preload it. But this setting is only a hint—iOS will never preload data because users could be on expensive mobile data networks. (Browser support for `preload` is relatively new.)

You may have also noticed the `type` attribute on the `<source>` elements, e.g.:

```
<source src="myaudio.mp3" type="audio/mpeg">
```

This attribute tells the browser what container format was used, so it can work out whether it supports the file format without having to start downloading it to check. Sadly, audio format support in modern browsers is a bit of a mess, as we'll see shortly.

This is just a brief description of the `<audio>` element. For *implementation* I suggest you use a HTML5-friendly, JavaScript-based media player. It will help smooth out the implementation issues and save you re-inventing the wheel, which is especially helpful given the immature `<audio>` implementations in current browsers. We'll look at our media player options at the end of this chapter. (But if you want something simple you can drop in right now, try `audio.js`: <http://kolber.github.com/audiojs/>.)

However, for audio file *preparation* you need to understand the issues around codecs which we'll discuss after we look at the `<video>` element. We'll also look at HTML5 audio for games later in the chapter too, and touch on the flaws and future of `<audio>`.

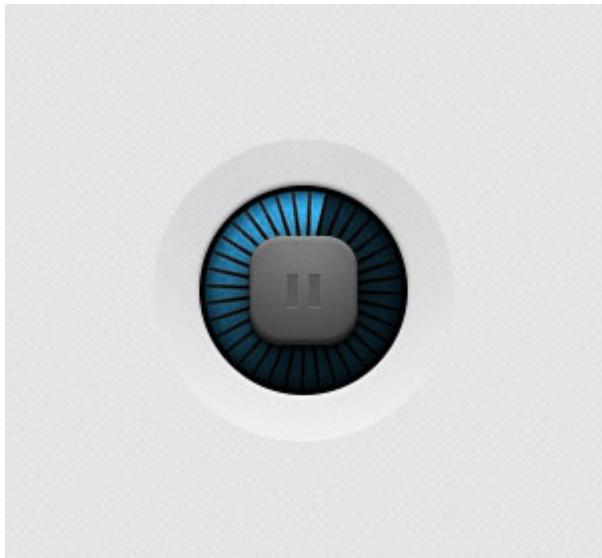
There is a JavaScript API for manipulating `<audio>` and `<video>` (allowing you to roll your own controls for playing, pausing, adjusting volume, etc.), which is covered in the resources and tutorials below.

For more `<audio>` resources, see:

- For the basics of the `<audio>` and `<video>` JavaScript API, see: <https://developer.mozilla.org/en/DOM/HTMLMediaElement>.
- Google's HTML5 Rocks "Quick Guide to Implementing the HTML5 Audio Tag (with Fallback to Flash)": <http://www.html5rocks.com/tutorials/audio/quick/>.
- Dev.Opera's "An HTML5 `<audio>` radio player": <http://dev.opera.com/articles/view/html5-audio-radio-player/>.
- "Building a Custom HTML5 Audio Player with jQuery" by Neutron Creations: <http://neutroncreations.com/blog/building-a-custom-html5-audio-player-with-jquery/>.

- For Safari and iOS implementations, see the Safari Developer Library: [http://developer.apple.com/library/safari/#documentation/AudioVideo/Conceptual/Using\\_HTML5\\_Audio\\_Video/Introduction/Introduction.html](http://developer.apple.com/library/safari/#documentation/AudioVideo/Conceptual/Using_HTML5_Audio_Video/Introduction/Introduction.html).
- SoundManager 2 is currently the go-to open source audio library for “reliable cross-platform audio under a single JavaScript API”: <http://www.schillmania.com/projects/soundmanager2/>.

The beauty of having the `<audio>` and `<video>` elements available as ordinary HTML is you can style them with CSS, including advanced CSS3. Check out the beautiful Zen Audio Player (playing Girl Talk no less) to see what’s possible: <http://lab.simurai.com/zen-player/>.



*Figure 10.2. The Zen Audio Player really is a thing of beauty—be sure to see it in action.*

The crucial thing to understand with `<audio>` is the codec situation, so keep reading and we’ll get delve into it after a brief trip through the `<video>` element.

## The `<video>` Element

For video, we can use:

```

<video controls autoplay loop muted preload="auto"
poster="myvideobackground.jpg" height="250" width="300">
  <source src="myvideo.webm" type="video/webm">
  <source src="myvideo.mp4" type="video/mp4">
  <!-- Fallback content, such as a Flash player, or a link
to the file here -->
</video>

```

(If you want to see what’s *really* involved with `<video>` implementation, including browser issues, setting MIME types and more, see Kroc Camen’s wonderfully thorough “Video for Everybody” article: [http://camendesign.com/code/video\\_for\\_everybody](http://camendesign.com/code/video_for_everybody).)

As you can see, it’s a similar setup to the `<audio>` example. In fact, the `controls`, `autoplay`, `loop`, `muted` and `preload` attributes all behave the same way. But they’re not *quite* the devil’s tools as they are for audio. Here we could have an autoplay, looping video advertisement that’s muted until the user decides otherwise.

And just like `<audio>` we have to deal with codec support issues by specifying multiple video files for maximum HTML5 compatibility. We use the `<source>` tag to give browsers a list of video files, and they use either the first one they support or the fallback content (such as a Flash player). The `type` attribute gives browsers a hint as to which file they should try and play. We’ll discuss this after looking at the codec situation.

The `<video>` element has several of its own unique attributes. The main one is `poster`, which is the static image that’s displayed until the first frame of the video is available. This may only be for a second or two in some cases, but on mobile devices (such as iOS) the poster is shown until the user initiates playback.

(At least that’s the theory. IE9’s handling of the `poster` image is quirky, as Ian Devlin found: <http://www.iandevlin.com/blog/2011/12/html5/the-problem-with-the-poster-attribute>. A bug in iOS 3 prevents video playback when `poster` and `<source>` elements are used. An iPad running iOS 3.x ignores everything but

the first `<source>` element. And Android 2.x's entire `<video>` implementation seems to be one giant bug. See the Video for Everybody article above for more.)

The `height` and `width` attributes are also specific to the `<video>` element. But again, because the `<video>` element is just another bit of HTML it can be styled and manipulated with CSS, including advanced CSS3. You can transform and animate the video itself, add shadows, and so on. This is one of the coolest things about `<video>` being just another element in the DOM. You can even use the `<canvas>` element to manipulate your video source, as discussed here:

<http://html5doctor.com/video-canvas-magic/>.

## Video Accessibility

Media accessibility is also being developed. A `<track>` element has been added to the spec to provide captioning. Or, as the spec puts it: “The track element allows authors to specify explicit external timed text tracks for media elements” (<http://dev.w3.org/html5/spec/Overview.html#the-track-element>). The `<track>` element sits between the `<video></video>` tags, and looks like this:

```
<track kind="subtitles" src="moviecaptions.en.vtt"
srclang="en" label="English">
```

You can find presentations and discussions on the issues and proposed solutions here: <http://blog.gingertech.net/2011/03/29/webvtt-explained/> and here: <http://www.iandevlin.com/blog/2011/05/html5/webvtt-and-video-subtitles>.

Currently only IE10 and Chrome 18 (in beta as I write) support `<track>`. For more, see “Getting started with the HTML5 track element” (<http://www.html5rocks.com/en/tutorials/track/basics/>).

## API And Resources

The new HTML5 JavaScript API for media also handles video playback, which lets you roll your own controls. It's covered in the resources and tutorials below.

For more `<video>` resources, see:

- “Video for Everybody” is worth a read and/or bookmark: [http://camendesign.com/code/video\\_for\\_everybody](http://camendesign.com/code/video_for_everybody)

- Google’s HTML5 Rocks HTML5 video article covers the basics and some pretty out-there examples, including SVG and video:  
<http://www.html5rocks.com/en/tutorials/video/basics/>.
- Mark Pilgrim’s Dive Into HTML5 chapter on video has all the gory details on codecs (and the difference between codecs and container files), encoding, browser support, server MIME types and more:  
<http://diveintohtml5.info/video.html>.
- Dev.Opera has an extensive guide on HTML5 video, optimistically titled “Everything you need to know about HTML5 video and audio”:  
<http://dev.opera.com/articles/view/everything-you-need-to-know-about-html5-video-and-audio/>.
- The Safari Developer Library guide to HTML5 audio and video has all the ins and outs for Safari and iOS related implementation:  
[http://developer.apple.com/library/safari/#documentation/AudioVideo/Conceptual/Using\\_HTML5\\_Audio\\_Video/Introduction/Introduction.html](http://developer.apple.com/library/safari/#documentation/AudioVideo/Conceptual/Using_HTML5_Audio_Video/Introduction/Introduction.html).

The savior for `<video>` is standalone JavaScript players, which let us use one file, HTML5 (and a given codec) where it’s supported, and Flash everywhere else.

We’ll get to the media players in a moment. In the meantime, grab a fistful of hair, and get ready to pull.

## Codex, You’re Killing Me

Okay, so HTML5 for modern browsers, and Flash as a fallback for older browsers. Got it.

Not so fast. This is HTML, which isn’t so much “Anything that can go wrong, will go wrong” as “Anything that can cause disagreement, will cause disagreement”. And the disagreement here is codecs, for both audio and video.

If you use an image tag, all browsers can display the image whether it’s a JPEG, GIF or PNG—there’s no mandatory format.

The HTML5 spec (reluctantly) takes a similar view with the `<audio>` and `<video>` tags by not specifying a specific format (i.e. codec) for audio or video.

You specify the format you want to use, and it's up to the browser to support it (or not, as we'll see).

Now if browser vendors all agreed on a single format (or several formats), we'd have universal HTML5 audio and video in all modern browsers. Unfortunately, that hasn't happened. Here's the HTML5 editor Ian Hickson on the situation in mid-2009 (<http://lists.whatwg.org/pipermail/whatwg-whatwg.org/2009-June/020620.html>):

*After an inordinate amount of discussions, both in public and privately, on the situation regarding codecs for <video> and <audio> in HTML5, I have reluctantly come to the conclusion that there is no suitable codec that all vendors are willing to implement and ship.*

*I have therefore removed the two subsections in the HTML5 spec in which codecs would have been required, and have instead left the matter undefined, as has in the past been done with other features like <img> and image formats, <embed> and plugin APIs, or Web fonts and font formats.*

## The Patent Problem

The problem with agreeing on a codec comes down to patents. Some media formats—including MP3 for audio (yes, the humble .mp3), and the popular H.264 format for video (usually used in .mp4 and .mkv files)—have patents that make companies cough up licensing fees to use the decoders in their products.

For big companies like Apple, Microsoft and Adobe (with Flash) this isn't a problem—they support both MP3 and H.264. But for ideological and financial reasons Opera and Mozilla don't support MP3 for audio, or H.264 for video. (In early 2011 Google threatened to drop H.264 support from Chrome on the desktop, but as of early 2012 that has yet to happen. See the post: <http://blog.chromium.org/2011/01/html-video-codec-support-in-chrome.html>.)

Yep, it's a format war. And one that isn't likely to be resolved any time soon.

What are the alternatives? In the audio department, Mozilla and Opera champion the patent-free Ogg format (ditto for video, but it's seen as inferior). And in mid-2010 Google released the theoretically patent-free WebM video format (after buying it for a cool \$100m) to provide a "Can't we all just get along?" solution for video that everyone could use, thereby resolving the deadlock.

So everyone could just switch to those, right? Not quite. For something to be truly "patent-free", it generally has to be proven in court. So Microsoft and Apple take a "better the devil you know" approach and pay royalties for the codecs they use, especially with video and H.264. They figure it's better to do that than opt into supposedly "patent-free" technology that may not be so patent-free after all and could make them liable in the future. Indeed, questions are already being raised about potential WebM patent infringement, so these are valid concerns. (That's the condensed version, in any case!)

## H.264 Is Baked In

Even if everyone *could* use WebM for video, it's not like Apple *or* Google (with Android) could just push out a software update and have everyone running WebM for video (especially on mobile).

Why not?

H.264 uses hardware acceleration in mobile devices (as well as desktop and other devices), which is how we get to watch high-quality video on low-powered devices without destroying battery life.

Throw in additional issues, such as the industry toolchain built around H.264, and the situation gets even murkier. You can see why a wholesale switch from H.264 would be difficult (at least in the short- to medium-term) even if Apple, Microsoft and others *did* decide to move in that direction. As one Hacker News commenter put it (<http://news.ycombinator.com/item?id=2106285>):

*The digital video world runs on H.264, it has deep, complicated, expensive internal toolchains to support it, legacy archives encoded in it, basically entire businesses built around it. The video production world is far larger and more complex than you're picturing it.*

H.264 on mobile is, for all vendors, and for the foreseeable future, a fact of life.

## Google Threatens To Take Chrome WebM Only, And Then Doesn't

Let's look at who supports what.

Google, as mentioned, announced in early 2011 they were going to *remove* H.264 support from Chrome to focus on WebM (see the announcement: <http://blog.chromium.org/2011/01/html-video-codec-support-in-chrome.html>), but as of writing in early 2012 that has yet to materialize. Google-owned YouTube is transcoding all their videos into WebM so they can serve both WebM and H.264. Yep, they're (supposedly) duplicating their entire YouTube library. And in April 2011 they announced they had "already transcoded videos that make up 99% of views on the site or nearly 30% of all videos" (<http://youtube-global.blogspot.com/2011/04/mmm-mmm-good-youtube-videos-now-served.html>).

(Fun fact: they also mentioned that, as of mid-2011, six *years* of video was uploaded to YouTube *every day*.)

As for Microsoft, they share Apple's position and are sticking with H.264 (see: <http://www.fastcompany.com/1723373/microsoft-sides-with-apple-over-google-on-h264-video>). They allow native WebM playback *if* a user installs the codec, but there's no shipping support.

Google and Microsoft then released tit-for-tat plugins. Google released a WebM plugin for IE9 (<http://tools.google.com/dlpage/webmmf>), though whether it gets any significant adoption remains to be seen. Microsoft has in turn released a H.264 plugin for Chrome (see: <http://blogs.msdn.com/b/interoperability/archive/2011/02/01/greater-interoperability-for-windows-customers-with-html5-video.aspx>) and Firefox (see: <http://blogs.msdn.com/b/interoperability/archive/2010/12/15/html5-video-and-interop-firefox-add-on-provides-h-264-support-on-windows.aspx>) for its Windows 7 users that use Firefox.

Opera and Mozilla refuse to support H.264, and have added WebM support in Opera 10.6+, and Firefox 4+ respectively. (Firefox 3.x and earlier Opera

versions only offer Ogg Theora support.) Mozilla's position on H.264 appears to be shifting though, as we'll see below.

Adobe has said they'll support WebM in future versions of Flash, but it's unclear when full support will appear. (See the brief mention in this story: [http://news.cnet.com/8301-30685\\_3-20061315-264.html](http://news.cnet.com/8301-30685_3-20061315-264.html).) Flash 11 shipped in late 2011 without WebM support.

Even if a version of Flash does support WebM in the future, it will count for little. Apple doesn't support WebM in Safari, and WebM and Flash are unsupported on iOS; future Android devices won't support Flash (given Adobe has abandoned the Flash plugin on mobile) and all current Android devices don't support WebM; and Microsoft won't support Flash on their Windows Phone platform (or indeed on Metro), which means H.264 will be a necessity for video on mobile (at the very least) for legacy and future device support.

For a further breakdown on browser and device compatibility, see the handy chart here: <http://mediaelementjs.com>.

## Codecs: What To Do?

Phew! So where does that leave us?

If you want maximum *native* HTML5 support—including Firefox and Opera—you need to store two copies of your audio (MP3 and Ogg Vorbis), and potentially three copies of your video files (H.264, WebM, and Ogg Theora for legacy support in Firefox 3.x). Having to encode and store so many different versions is a royal pain in the ass, but there you go.

Alternatively, you can:

- Use MP3 for audio, which works natively in Safari, IE9, and Chrome (including iOS and Android).
- Use H.264 for video, which works natively in iOS, Safari and IE9.
- Use a media player (or script your own) that uses the Flash player for older and non-codec-supporting devices, given Flash can play MP3 and H.264 in any Flash-enabled browser or device.

Given we need H.264 for iOS (and mobile in general), this scenario—which I would wager will be the most common—means Firefox, Opera, and Chrome (for video) will end up getting Flash. That’s right—those who stood up for free, open software will end up with proprietary, closed Flash. That is more ironic than a hipster’s mustache.

(Alternatively you can, as mentioned, double encode your video and provide a separate WebM file in a secondary `<source>` element for broad native HTML5 support in modern browsers.)

## Reality Bites

The reality may be a little different though. For one, Google has yet to drop H.264 support, and (as of writing) Mozilla is reconsidering their opposition to supporting the codec.

Andreas Gal, Director of Research for the Mozilla Corporation, kicked off the debate in March 2012 on a Mozilla mailing list, suggesting Mozilla (<https://groups.google.com/forum/#!msg/mozilla.dev.platform/-xTei5rYThU/DkM9A1bkNNIJ>):

*[S]upport decoding any video/audio format that is supported by existing decoders present on the system, including H.264 and MP3. There is really no justification to stop our users from using system decoders already on the device, so we will not filter any formats.*

This would mean Mozilla could ship Firefox (for mobile platforms, particularly) with support for the *operating system*’s licensed decoders (where available); avoid the millions of dollars in licensing fees; and maintain at least the veneer of ideological purity. This is actually about much more than Firefox—Mozilla are working on their own web-based mobile *platform* called Boot to Gecko (or B2G, which we’ll look at in chapter twelve), but it’s not going to have much of a chance if it can’t play H.264 video (and MP3 audio) like the other major platforms. (Note that Gal’s position in the Mozilla Corporation doesn’t give him any particular sway over the direction of Firefox or the Mozilla Project, he’s just acting as a contributor.)

But wasn't Google supposed to have solved this problem with WebM? Here's Gal again (<https://groups.google.com/forum/#!msg/mozilla.dev.platform/-xTei5rYThU/iZ767IwV1jUJ>):

*Google pledged many things they didn't follow through with and our users and our project are paying the price. H.264 wont go away. Holding out just a little longer buys us exactly nothing.*

Mozilla is in an awkward position, as Firefox developer Justin Dolske noted in the mailing list discussion ([https://groups.google.com/forum/#!msg/mozilla.dev.platform/-xTei5rYThU/3d6e-Sgo\\_ZQJ](https://groups.google.com/forum/#!msg/mozilla.dev.platform/-xTei5rYThU/3d6e-Sgo_ZQJ)):

*But I think if Mozilla is going to do an about-face on open video standards (and it is an about-face), then there should be some serious discussion about it. Certainly more than than a few terse words saying it's hopeless and obvious. That makes it sound more like a half-hearted notification of a decision that's already final.*

*At the very least, it needs to be explained enough so the community can understand the change. We spent a lot of time and made a lot of blog posts about why H.264 was bad for the web. Leaving those who advocated for us suddenly high-and-dry doesn't feel like the right thing to do.*

H.264 has the momentum, and Mozilla face an unenviable choice between pragmatism and ideology. H.264 is, despite its non-free nature, a de facto standard (much like MP3) and they will need to support it or be left behind. (For more, see Ars Technica's report: <http://arstechnica.com/gadgets/news/2012/03/idealism-vs-pragmatism-mozilla-debates-supporting-h264-video-playback.ars>. Keep an eye on Ars' coverage for updates as they develop.)

And making it more or less official, Mozilla's CTO, Brendan Eich, threw his weight behind Gal's proposal in "Video, Mobile, and the Open Web" (<http://hacks.mozilla.org/2012/03/video-mobile-and-the-open-web/>), writing:

*What I do know for certain is this: H.264 is absolutely required right now to compete on mobile. I do not believe that we can reject H.264*

*content in Firefox on Android or in B2G and survive the shift to mobile.*

*Losing a battle is a bitter experience. I won't sugar-coat this pill. But we must swallow it if we are to succeed in our mobile initiatives. Failure on mobile is too likely to consign Mozilla to decline and irrelevance. So I am fully in favor of Andreas's proposal.*

If this is the route Mozilla take, it's difficult to imagine Opera holding out indefinitely. The practical implication for this is encoding media in H.264 (with a Flash fallback for older browsers) and MP3 (ditto) will be sufficient in the long term.

## Video Types... Oh My

Now you (hopefully) understand the complexity of the codec situation, and the way browser support for each codec differs, let's look at how we can tell browsers which codecs we're using, so they can make smart decisions about which video they load.

That brings us to the one video attribute we haven't discussed—the `type` attribute on the `<source>` element, e.g.:

```
<source src="myvideo.mp4" type="video/mp4">
```

This attribute tells the browser what *container* and *codec* is used for the video specified in the `src` attribute. In the above example we've only specified what *container* format is used by listing its MIME type (i.e. media format type). The MIME type above tells the browser “This file is a video using the mp4 container format”. *Container formats* like mp4 are a little like zip files, in that they are simply a container for the *actual* video and audio files, which are encoded with specific *codecs*, and wrapped up to make the final video file. (The `type` attribute can also be used on the `<video>` element itself, not just a nested `<source>` element, if you're only using one file. Same goes for `<audio>`.)

The information we put in the `type` attribute is just a hint to the browser, but it's not necessary for browsers to play the video. What *is* necessary is to edit your

.htaccess file to make sure your server sends this files with the right MIME type, as the instructions here describe: <http://mediaelementjs.com/>, and are covered elsewhere (e.g. the “Video for Everybody” article mentioned earlier).

We can also specify both the container format *and* codec used, e.g.:

```
<source src="myvideo.mp4" type='video/mp4;
codecs="avc1.42E01E, mp4a.40.2"'>
```

Here we’ve specified the container format, *and* the codecs for both the video and the audio in the source file. (The video codec is a flavor of H.264, along with the AAC audio codec. Note we also have to use single quotes for the `type` attribute here, as the codecs parameter uses double quotes.)

What’s the point of all this? Well, specifying the `type` attribute means the browser doesn’t have to start downloading every listed file just to check if it can play it or not. It can scan the markup, and potentially start preloading the video file it supports. Leaving it out, however, isn’t the end of the world. From “The State Of HTML5 Video” by LongTail Video, in 2012

(<http://www.longtailvideo.com/html5/>):

*Every browser supports the <source> tag for loading multiple sources. Our tests show that including the type attribute prevents some preloading, but breaks compatibility with Android 2.2. Setting the codecs in the type attribute has no impact in any browser.*

Note the Android 2.2 compatibility issue. The comment about codecs suggests that from LongTail’s testing we (by and large) only need to specify the *container* and the browser will have a go at loading it.

### Querying Supported Video Types With JavaScript

We can also query the browser using the `<video>` JavaScript API and its `canPlayType()` method to see which formats the browser supports. For example, with the codecs we specified above (“avc1.42E01E, mp4a.40.2”), browsers that support these formats (Safari and IE9+) will respond with `probably`, which is as close to “Yes, we support this file” as we get in HTML5.

If we only specify the container format ('video/mp4'), Safari and IE9+ (for example) respond with *maybe*, because they know they can read that container format, but don't know what codecs lie inside. Browsers that don't support a given container or codec format just return a null string.

Things get awfully complicated though. Here are the three variables we have to deal with:

- **Browser response:** Due to the complexities of encoding media (and especially video), browsers only know for sure that they *can't* play formats they don't understand. In those cases, they return a null string (when queried through the `canPlayType()` method). Beyond that, the HTML5 spec says they must return either *maybe* or *probably* depending on the browser's confidence it can play a certain file based on the information we've provided.
- **Containers and codecs:** There are *container* formats such as mp4, and there are the actual codecs, or more accurately, *flavors* of the actual codecs (such as avc1.42E01E), which can be queried.
- **Browser support:** Finally, as we've seen, the codec support from the major browsers is a pretty complex situation.

Therefore, we have multiple browsers, supporting multiple container/codec varieties (for both audio and video), and giving one of three responses.

Thankfully, the WHATWG maintain a table of browser responses so we can see which response for a given container format, or container and codec combination, we *should* get from a given browser: [http://wiki.whatwg.org/wiki/Video\\_type\\_parameters#Browser\\_Support](http://wiki.whatwg.org/wiki/Video_type_parameters#Browser_Support). Microsoft have a small script demonstrating how this works here: [http://msdn.microsoft.com/de-de/library/hh325437\(v=vs.85\).aspx](http://msdn.microsoft.com/de-de/library/hh325437(v=vs.85).aspx).

(Also note that browser responses we *actually* get can be buggy, as this slightly dated post from mid-2010 suggests: <http://rakaz.nl/2010/06/problems-with-html5-video-codec-detection.html>.)

## Audio And Video Media Players To The Rescue

What a mess.

Fortunately, people have written tools to take the pain out of serving the right video (or audio) to the right browser. These media players hold your hand through the whole codecs mess and legacy support issues; provide plenty of customization options; and generally smooth out the whole implementation process.

Here are a few examples:

## MediaElement (Video And Audio, Free)

<http://mediaelementjs.com>

MediaElement.js

Home Blog How it Works Plugins Installation Options Support

Default Skin TED Skin WMP Skin

echo the life

<video src="echo\_conference" native

00:00 00:00

<audio src="Air\_Review" native

HTML5 <video> and <audio> made easy.  
One file. Any browser. Same UI.

- HTML5 audio and video players in pure HTML and CSS.
- Custom Flash and Silverlight players that mimic the [HTML5 MediaElement API](#) for older browsers
- Accessibility standards including [WebVTT](#)
- Plugins for [WordPress](#), [Drupal](#), [Joomla](#), [jQuery](#), and [BlogEngine.NET](#), [ruby\\_gems](#), [phone](#)
- See the [MediaElement.js Github page](#)
- Tell some folks:

Tweet 4,307 +1 558

Download Latest (2.7.0)

### What they're saying



I've been recommending MEjs to my audiences in my HTML5 talks recently. great stuff  
*Paul Irish*



Now that we have this, it's a done deal!  
*Chris Coyier*

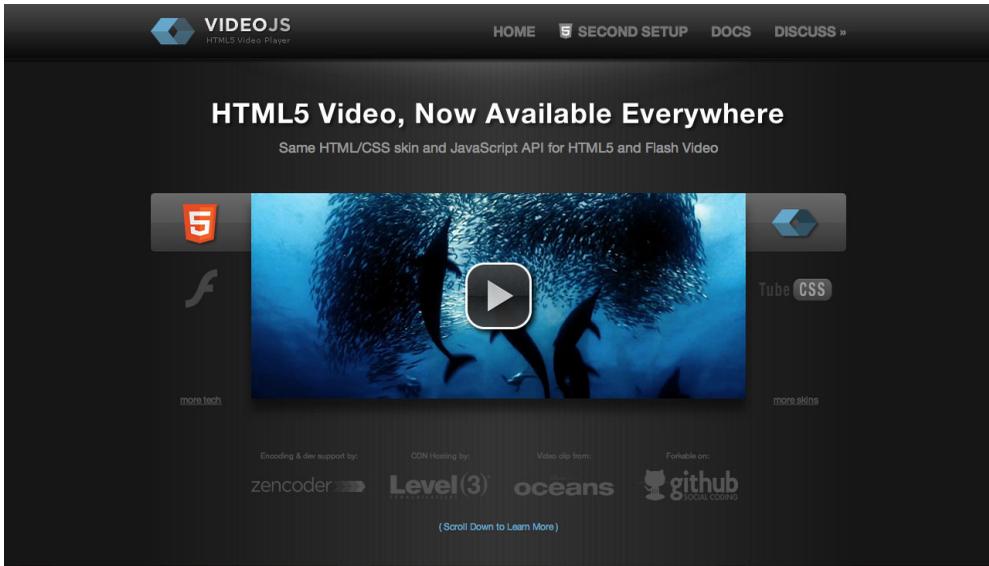


John and I have kids about the same age, and we both like JavaScript.  
*Nathan Smith*

A popular audio and video player that lets you use one file (a H.264 video file, for example) and deploy a consistent UI across all devices using Flash (or Silverlight) for playback where H.264 isn't natively supported. It comes with a jQuery plugin, as well as plugins for Drupal and Wordpress.

## VideoJS (Video, Free)

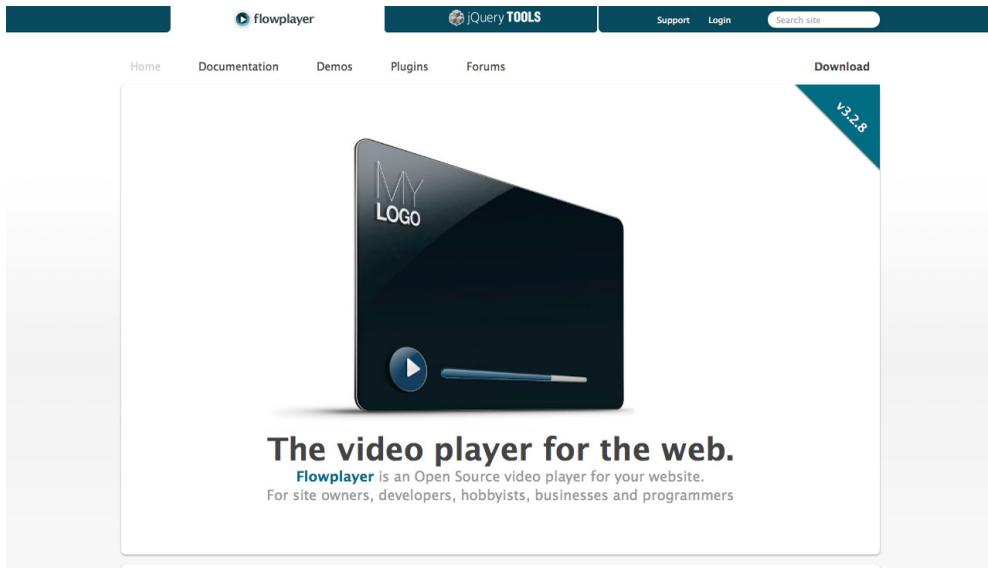
<http://videojs.com>



VideoJS is a slick HTML5 video player that offers some familiar, CSS-based skins, and similar broad support using Flash fallbacks for HTML5 video. It uses the markup from Video for Everyone (we linked to this earlier), and adds JavaScript for broader compatibility and CSS skinning options.

## Flowplayer (Video, Free And Commercial)

<http://flowplayer.org>



A free, open source player with Flowplayer branding. There's also a commercial offering (with support option) without branding.

## More Media Players

There's a variety of other players out there, including:

- **jPlayer** (<http://www.jplayer.org>), a free, open source, audio and video player (with handy playlist support)
- **Open Standard Media (OSM) Player** (<http://www.mediafront.org/project/osmplayer>), a free audio and video player written in jQuery that has a visual playlist
- **JW Player** (<http://www.longtailvideo.com/players/>), another option with HTML5 support
- **SublimeVideo** (<http://sublimevideo.net/>), a paid, hosted, HTML5 player solution.
- **Popcorn.js** (<http://popcornjs.org/>) is a “JavaScript event framework” and is part of the Mozilla Popcorn project: <http://mozillapopcorn.org/>. It's particularly useful if you want to trigger updates of other content on the page in sync with the video.

(Of course, you can always use YouTube or Vimeo and let the native iOS playback take care of itself!)

With these available there's no need to roll your custom solution. Just grab one off the shelf and skin to your heart's content.

## Other Flies In The HTML5 Video Ointment: DRM, Streaming, And Full Screen Video

We've got basic `<video>` and `<audio>` embedding down, but even with our handy, ready-to-go media players and their Flash fallbacks there are still some significant missing (or very immature) features that HTML5 video lacks, and Flash supports. Keep in mind that Microsoft will *not* support Flash in IE10 in the Metro interface of Windows 8 (it will kick users out to the "Desktop" interface if they want to use Flash), so there will be considerable pressure to add these features sooner rather than later. When it comes to DRM, that may not be a good thing.

### DRM

In February 2012, representatives from Google, Microsoft and Netflix submitted the Encrypted Media Extensions v0.1 draft proposal (<http://dvcs.w3.org/hg/html-media/raw-file/tip/encrypted-media/encrypted-media.html>) to the W3C's HTML Working Group. The abstract for the proposal said:

*This proposal extends HTMLMediaElement to enable playback of protected content. The proposed API supports use cases ranging from simple clear key decryption to high value video (given an appropriate user agent implementation). License/key exchange is controlled by the application, facilitating the development of robust playback applications supporting a range of content decryption and protection technologies. No "DRM" is added to the HTML5 specification, and only simple clear key decryption is required as a common baseline.*

That is, this would provide a mechanism for DRM to happen on top of HTML5, for both `<audio>` and `<video>`, by extending their JavaScript API (the

`HTMLMediaElement` interface). The HTML5 editor, Ian Hickson, responded with (<http://lists.w3.org/Archives/Public/public-html/2012Feb/0274.html>):

*I believe this proposal is unethical and that we should not pursue it.*

Hickson also stated “DRM is evil” and then further elaborated on his reasons for rejecting the proposal outright here: [https://www.w3.org/Bugs/Public/show\\_bug.cgi?id=10902#c24](https://www.w3.org/Bugs/Public/show_bug.cgi?id=10902#c24). Mozilla also expressed concerns, given DRM and an open source browser are generally mutually exclusive (see this Ars Technica article for more: <http://arstechnica.com/business/news/2012/02/unethical-html-video-copy-protection-proposal-criticized-by-standards-stakeholders.ars>).

I mostly agree with Hickson. DRM is evil. Remember PlaysForSure? No? Exactly.

There is one important difference between DRM for video and DRM for audio though, and that’s streaming. We have a culture of streaming (and renting) video, in a way that we didn’t with audio when the music DRM wars were raging. DRM sucked when it was applied to music you bought, because if the DRM platform died (and they did), so did your music collection. The temporal nature of streaming media alleviates these concerns, but only to a point. (And if it can be implemented for streaming media, it’s not hard to imagine the media majors insisting on DRM for purchased content, too.)

These issues around streaming “protected” content explains Netflix’s interest in seeing some sort of DRM available when using web standards. Netflix is keen to use HTML5 (in the broad “web platform” sense) where they can (as they discuss here: <http://techblog.netflix.com/2010/12/why-we-choose-html5-for-user.html>), but they obviously feel they need some kind of DRM system in place to stream the content they license with HTML5. (They also need an actual streaming protocol, as we’ll see next.) Some at Google and Microsoft evidently feel this is a necessity too. (Note though that Hickson also works for Google, so it’s not a company-wide position.)

Microsoft is in a particularly interesting position, given they won’t support Flash in IE10 in Metro, but will obviously want to provide some mechanism so their users can still access “protected” streaming video services within IE10 on Metro.

The alternative, of course, is the streaming video companies instead implement their services as Silverlight-powered Metro apps, and move off the web altogether for those users.

This may be another case where the post-Flash web results not in an open standards utopia, but in a return to platform-specific apps. That said, given the cost of implementing DRM, that may be a price the standards movement is willing to bear.

As I write, debate rages on the W3C mailing lists, with a summary here:

<http://lists.w3.org/Archives/Public/public-html/2012Mar/0087.html>.

But that debate may be moot. In March 2012, Philippe Le Hegaret from the W3C wrote (<http://lists.w3.org/Archives/Public/public-html/2012Mar/0097.html>):

*[L]et's be clear: W3C has many participants interested in finding a solution around media content protection. So, we are definitively interested in the space, independently of whether the HTML Working Group is interested in developing a solution or whether it is done in a separate group. Whatever we choose, we will do our best to get the right balance between producers and consumers.*

That is, we're going to implement DRM with or without you, because our paid-up members want it. Ominous stuff.

## Streaming

We've touched on the DRM aspects of streaming, but what about the other technical challenges? In late 2010 Netflix discussed a number of issues with streaming and HTML5 they had identified, and have been busy working on: <http://techblog.netflix.com/2010/12/html5-and-video-streaming.html>. And in late 2011 Ars Technica posted an excellent article "The trials and tribulations of HTML video in the post-Flash era" (<http://arstechnica.com/business/news/2011/11/the-trials-and-tribulations-of-html-video-in-the-post-flash-era.ars>) which elaborates on some of those issues, including streaming. The Ars article also notes:

*[T]ransitioning video delivery in the browser from Flash to HTML5 will also pose some major challenges for content creators. The*

*standards aren't fully mature yet and there are still a number of features that aren't supported or widely available across browsers.*

*For an illustration of how deep the problems run, you need only look at Mozilla's Firefox Live promotional website, which touts the organization's commitment to the open Web and shows live streaming videos of Red Panda cubs from the Knoxville Zoo. The video is streamed with Flash instead of using standards-based open Web technologies.*

*In an FAQ attached to the site, Mozilla says that it simply couldn't find a high-volume live streaming solution based on open codecs and open standards. If Mozilla can't figure out how to stream its cuddly mascot with open standards, it means there is still work to do.*

A streaming standard has been in the works for some time: Dynamic Adaptive Streaming over HTTP (DASH), which has the support of Microsoft, but there's still a ways to go there before support materializes. DASH is also codec agnostic—it doesn't resolve the codec impasse we discussed earlier. (For more see “What is MPEG DASH?”: <http://www.streamingmedia.com/Articles/ReadArticle.aspx?ArticleID=79041>.)

Apple currently has its own streaming protocol, HTTP Live Streaming (HLS), which is used to deliver content to its iOS devices. Google added support in Android 3.0+, and it allows encrypted data and works with third-party DRM solutions. (See “What is HLS (HTTP Live Streaming)?” for more: <http://www.streamingmedia.com/Articles/Editorial/What-Is-.../What-is-HLS-%28HTTP-Live-Streaming%29-78221.aspx>.)

DRM, the codec impasse, support for a new technical standard, and competing existing standards. That's a lot of issues that need to be resolved before standards-based streaming becomes a reality.

## Full Screen

Finally, one of the most common things we do with Flash video is make it full screen. That's not possible in HTML5. There is, however, a Fullscreen API W3C specification (<http://dvcs.w3.org/hg/fullscreen/raw-file/tip/Overview.html>) that

has experimental support in Firefox 10+ and recent WebKit browsers (Chrome 15+ and Safari 5.1+). It's unclear whether IE10 will support this feature. Interestingly, the Fullscreen API can make *any* element full screen, including (for example) the `<canvas>` element, and it could be used for a full screen reading mode too, for example.

For more on the Fullscreen API, see:

- The Mozilla Hacks blog has a tutorial (including styling information) and a demo you can play with. Tutorial: <http://hacks.mozilla.org/2012/01/using-the-fullscreen-api-in-web-browsers/> and demo: <http://robnyman.github.com/fullscreen/>.
- There's another tutorial here: <http://tutorialzine.com/2012/02/enhance-your-website-fullscreen-api/>.
- Keep an eye on <http://caniuse.com/#feat=fullscreen> for browser support stats and more resources.

## Is HTML5 `<audio>` Ready For Games?

One final note on `<audio>` and its potential for gaming.

Dominic Szablewski posted an epic (and profanity-laced!) rant on the state of HTML5 audio, in relation to developing HTML5 games. It's well worth a read for entertainment value, and gives some indication of the (im)maturity of HTML5 audio support, particularly for interactive purposes. Szablewski says.

*Surprisingly, Google's Chrome has the worst HTML5 Audio support of all the good Desktop Browsers - that is every Browser but IE. I'm not an audio engineer, but before Browser vendors took their shot at it, my impression of digital audio was that it is a solved problem. I'm amazed that after so many iterations HTML5 Audio is still that broken.*

*The Audio support on mobile Browsers (iOS and Android) is laughable at best. It's completely unusable for even the simplest of tasks. You can jump through hoops and ask real nice, but it still sucks ass.*

(See the article for more: <http://www.phoboslab.org/log/2011/03/the-state-of-html5-audio>. I love the Scumbag Steve hat.)

The short answer? No, it's not really ready for games.

In fact, it's debatable whether an advanced audio API will make the HTML spec anytime soon, given Ian Hickson's feelings on the matter. In June 2011 he wrote (<http://lists.w3.org/Archives/Public/public-audio/2011AprJun/0118.html>):

*I'm not convinced that audio APIs are really the next most important thing to work on on the Web platform. [...] [M]aybe the best Audio API for the Web platform for the near future is no API at all.*

In the meantime, Flash will remain a necessity for advanced, cross-browser audio support for games, as the team at Pixel Lab found with their excellent Canvas-powered version of Cut the Rope (play it here: <http://www.cuttherope.ie/>)

Robby Ingebretsen described how hard the Pixel Lab team tried to get HTML5 audio happening for all browsers, but (<http://nerdplusart.com/why-is-there-flash-in-the-html5-version-of-cut-the-rope/>):

*We were up against browser quirks and bugs, not just feature support. In other words, even if a browser supported HTML5 audio we weren't guaranteed that it would reliably handle the complex sound requirements of the game.*

Microsoft sponsored the port of the game to HTML5 to show off IE9's standards support, and fortunately for Ingebretsen they got audio working in IE9 with `<audio>` alone, as apparently IE9's `<audio>` implementation is quite good. In other browsers, however, the bugs proved difficult to overcome. Therefore, they chose to use the "very capable" SoundManager 2 library, which comes with a Flash fallback they could rely on instead. The SoundManager 2 website describes the project (<http://www.schillmania.com/projects/soundmanager2/>):

*SoundManager 2 gives you a single, powerful API that supports both new and old, using HTML5 audio where supported and optional Flash-based fallback where needed. Ideally when using SoundManager 2, audio "just works."*

The Flash fallback meant users with browsers that have a flakey `<audio>` implementation still got a great audio experience when playing the game. It just wasn't pure HTML5, but hats off to the team for a great port and doing everything they could to get to 100% HTML5 support (and getting there in IE9). Their experience does, however, highlight the weakness of `<audio>` implementations for HTML5 games.

## The Future Of Audio

That's not to say things are sitting still with `<audio>` in the browser, or that an audio API must be a part of HTML5.

Some clever workarounds have appeared, such as Remy Sharp's audio sprites (<http://remysharp.com/2010/12/23/audio-sprites/>), and the sophisticated Web Audio API from Google is currently a W3C proposal, and ships in Chrome. Firefox 4+ has its own Audio Data API, and both API's are used in emerging JavaScript audio libraries (see: [https://wiki.mozilla.org/Audio\\_Data\\_API#JavaScript\\_Audio\\_Libraries](https://wiki.mozilla.org/Audio_Data_API#JavaScript_Audio_Libraries)). There is currently no Safari, IE, or Opera support for these APIs, however.

While it's exciting to see this innovation happening (albeit outside the HTML5 spec); broad, mature, consistent audio support for applications like gaming is still looking a long way off yet.

## Wrapping Up

Native support for `<audio>` and `<video>` is welcome, and necessary for mobile devices in particular. But keep in mind that the technology is still somewhat immature (especially on Android).

The codec issues will not be resolved any time soon, and iOS will still need H.264 for video. So we need to tread carefully. Don't assume it will work flawlessly just because it's in the spec.

My advice? Keep an eye on what your favorite media player supports, and let it do the hard work.



# THE TRUTH ABOUT SVG: THE FLASH CHALLENGER THAT WAS, WASN'T, & NOW...

We touched on the idea of Flash-challenging technology when we looked at Canvas in chapter nine. But perhaps the most serious challenger to Flash in the past decade has been Scalable Vector Graphics (SVG), an XML format for 2D graphics and animation, which is making yet another mini-comeback. Let's take a very brief look at the dead-again, less-dead-again technology that is SVG.

### SVG, SVG...

Oh, SVG (Scalable Vector Graphics). What can we say about you? You're a separate W3C spec that's been in development since 1998. You're not part of the HTML5 spec, but HTML5 *will* let you appear inline with other markup. You're all about vector shapes, which makes you the Illustrator to Canvas' Photoshop. You've promised so much, for so long—in 2002 people were writing 1000+ page books about you (SVG Unleashed, 2002, Sams). Yet you've never made it to the big leagues. What happened?

SVG is both a relic of past web standards and a technology that's finally kinda-sorta-almost arrived. It's an XML format for vector graphics (think of SVG being to graphics as HTML is to text), which means it looks like a bunch of angle bracketed-tags and attributes. Remember when we looked at the history of HTML, and how it was going to be all XML? Well SVG is part of that vision—the vision that didn't work out.

Here's what basic SVG looks like:

---

```
<svg id="mysvgexample" height="200" width="300"  
xmlns="http://www.w3.org/2000/svg">  
  <rect id="myrectangle" width="200" height="100"  
fill="red" x="20" y="20" />  
</svg>
```

Which renders:

## Generic SVG



*Figure 11.1. Our exciting SVG demo.*

Yep, a red rectangle. I'll give you a moment to recover.

A quick explanation of what's going on here: we've set the `<svg>` element to be 200px high and 300px wide, which acts as a transparent container for the shapes or lines we want to render. (I've given it a border with CSS so you can see its boundaries.) We then create a rectangle with the `<rect />` element, give it a `height` and `width`, and an `x` and `y` offset from the container `<svg>` element. Add a `fill` (I've used the keyword `red`, but you could use a hex value too) and we're done.

In HTML5 you can drop this code right into your HTML document, and supporting browsers (IE9+, FF4+, Safari 5.1+, Chrome 7+, Opera 11.6+, iOS 5+, Android 3+) will render it as appropriate.

## SVG: Browser Support Arrives At Last

SVG is interesting again because browser support is finally arriving. Currently, all modern and not-so-modern browsers, including IE9 (but excluding, you guessed it, IE6-8), support SVG embedded using the `<embed>` or `<object>` elements. (Android 2.x being the sole hold-out.) The IE situation isn't quite as bad as it sounds though, as we'll see next.

Browsers are also starting to support SVG images in the `<img>` tag, and even as CSS backgrounds (specifically IE9+, FF4+, and all recent Safari, iOS, Chrome and Opera releases).

And there's even some support for applying advanced SVG features such as filters (e.g. gaussian blur), clipping and transforms to non-SVG objects. Firefox support for this is particularly good; see <http://people.mozilla.com/~prouget/demos/mashup/video.xhtml> in FF4+ for a rough demo. Apart from Firefox, support is patchy. As far as the Photoshop-like filters go, Chrome and Opera offer good support, but SVG filters will only arrive in version 6 of Safari, and there's currently no iOS 5 or Android 4 support. IE9 doesn't support them either, but IE10 offers hardware-accelerated SVG filter effects (see: <http://blogs.msdn.com/b/ie/archive/2011/10/14/svg-filter-effects-in-ie10.aspx>).

That's what's happening on the bleeding edge of SVG support (for the latest support stats see: <http://caniuse.com/#search=svg>). But what about real-world, cross-browser SVG we can use right now, in just about any browser?

## Yes, There Is Real World SVG We Can Use Right Now

Dmitry Baranovskiy's excellent Raphaël JavaScript library lets us do some simple, cool, Flash-like things with SVG. Under active development for several years, it offers browser support right down to IE6 (thanks to VML translation in

IE). You can check it out here: <http://raphaeljs.com/>. (And make sure you check out the sister graphing library, gRaphaël, as well: <http://g.raphaeljs.com/>.)

We'll look at what Raphaël (and SVG on its own) can do in a moment, but for now it's worth keeping in mind that, thanks to Raphaël, simple SVG (including animation) can be used with broad browser support.

Imagine if, all of a sudden, we had to support devices with an extremely wide array of resolution sizes and densities, some of which didn't support Flash. Wouldn't vector graphics—crisp at any resolution, and scalable to any size—make life a lot easier? Hold that thought...

## Many Faces Of SVG

Let's look at different faces of SVG. There's:

- SVG, the gigantic spec that's been kicking around (and growing) for a decade.
- Advanced SVG as it's being implemented in cutting-edge browsers.
- SVG as we can use it today in real world situations with tools like Raphaël and jQuery SVG (<http://keith-wood.name/svg.html>).

SVG can be (and is being) used right now for a variety of situations with solid cross-support, including iOS devices, so its time may have finally arrived. And just in time, considering Adobe's abandonment of Flash in mobile browsers, and Microsoft's refusal to allow plugins in IE10 in the default Metro interface in Windows 8. (For more on the death of Flash, see the discussion in chapter nine where we looked at Canvas and Flash.) And that's not to mention the pressing need for interface elements that can scale up from a phone, to an iPad "retina" display, to a 27" or 30" screen.

## SVG In The '00s—The Great Hope That Wasn't

SVG has been on quite a journey. In its full pomp it's extremely impressive, and currently does as much (if not more) than current CSS3 implementations. Full SVG support does everything from animation, to Photoshop-like filter effects we touched on earlier (see this example in Chrome, Opera or Firefox: [216](http://svg-</a></p>
</div>
<div data-bbox=)

[wow.org/filterEffects/chiseled.svg](http://www.wow.org/filterEffects/chiseled.svg)), to custom fonts, masking, video, and of course drawing vector shapes. It's all there (in the spec at least—browser support has lagged badly). So in a way it's a lot like Flash.

Well, basic SVG is like Flash a decade ago, but without the browser support or developer tools. In fact, before Adobe bought Macromedia (who developed Flash) Adobe was championing SVG as an open alternative. And in 2002 some 160 million people were using their SVG viewer plugin (since discontinued, see <http://www.xml.com/pub/a/2002/07/03/adobesvg.html>).

To give you a taste of the hype around SVG in the early '00s, here's a quote from a 2002 article on Digital Web titled "SVG: The New Flash" ([http://www.digital-web.com/articles/svg\\_the\\_new\\_flash/](http://www.digital-web.com/articles/svg_the_new_flash/)):

*SVG should soon be widespread, and its non-proprietary nature will help to hasten the progress. Flash will continue to be the dominant standard for quite some time because of its large client base. However, SVG is rising quickly. The distribution of the SVG plug-in via browser manufacturers will quickly increase the installed user base, just as it did for Flash. Future versions of various browsers will include SVG viewers as standard, and some already do.*

But SVG never really took off. (Let it be said its "non-proprietary nature" hasn't counted for much in the past decade.) It couldn't touch Flash's installed base, and never had a development tool as designer-friendly as Flash. And when Adobe bought Macromedia in 2005, as far as vector graphics on the web went it was Flash or bust.

Still, SVG never really died either. And with browser support now improving rapidly, and the writing on the wall for Flash, perhaps SVG is due for yet another comeback.

## SVG Browser Support: Android, What The Hell? Oh, And IE...

One sticking point at the moment is that Android 2.x doesn't offer even *basic* SVG browser support, despite Google pushing SVG elsewhere (see: <http://googlecode.blogspot.com/2009/10/svg-at-google-and-in-internet-explorer.html>). It's already in WebKit, the engine behind Android's browser, it was just intentionally left out of Android 2.x "to save space" (see the comment and discussion here: <http://code.google.com/p/android/issues/detail?id=1376#c4>). Go figure. (The browser in Android 3.0, the tablet version of Android, *does* support basic SVG, and SVG support finally arrived for mobile Android devices in the unified 4.0 release too.)

IE6-8 fortunately has a couple of libraries available that can try to translate SVG into something it can understand.

Raphaël, the JavaScript library for working with SVG we touched on earlier, falls back to IE's old VML (Vector Markup Language) for compatibility. (It's similar to the Canvas emulation we looked at in chapter nine, and has similar limitations, i.e. VML is *slow*.)

There's also SVG Web (<http://code.google.com/p/svgweb/>), which translates SVG to Flash for older browsers that don't support SVG, including IE. (Sadly, this still doesn't help us out with Android 2.x. See: [http://groups.google.com/group/svg-web/browse\\_thread/thread/77fb6970f5f01e97](http://groups.google.com/group/svg-web/browse_thread/thread/77fb6970f5f01e97).)

There's also canvg (<http://code.google.com/p/canvg/>), which renders SVG in Canvas, and provides some SVG support for Android as a stopgap measure until Android 4.0 becomes widespread. This approach is discussed here: [http://www.kendoui.com/blogs/teamblog/posts/12-02-17/using\\_svg\\_on\\_android\\_2\\_x\\_and\\_kendo\\_ui\\_dataviz.aspx](http://www.kendoui.com/blogs/teamblog/posts/12-02-17/using_svg_on_android_2_x_and_kendo_ui_dataviz.aspx). (Fabric.js may also help: <http://fabricjs.com/>.)

While these tools make basic SVG a reality for just about any browser, remember this is only basic SVG, not the crazy Photoshop-like filters that have become part of the SVG spec.

## SVG Demos: What Is It Good For?

Vector graphics are useful in lots of instances—maps, charts, illustrations, logos, visualizations, resolution-independent interfaces, and so on. And the success of Flash has certainly demonstrated the need for animated vector graphics is there. What may otherwise be delivered in Flash can potentially be done in SVG, and therefore be available to iOS users.

Let’s take a look at some SVG demos to see what it is capable of. We’ll look at a few general examples, and then a few real-world Raphaël examples.

### SVG Girl



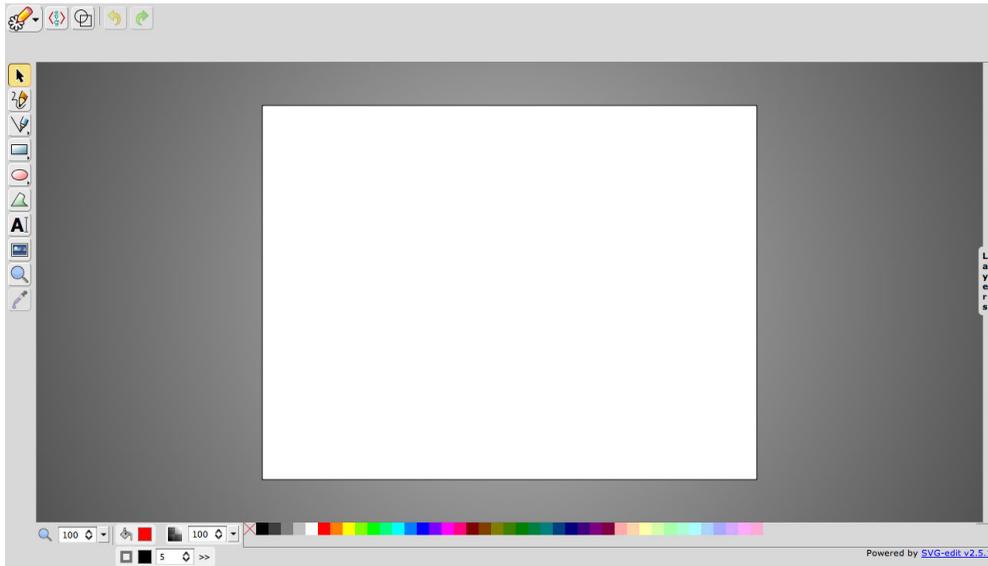
*Figure 11.2. The animated SVG Girl short is brief but impressive.*

SVG Girl is an “SVG animated video” Microsoft commissioned to show off their hardware accelerated SVG support in IE9 (though it works in any modern browser). It’s a brief but incredibly complex and impressive anime clip done with SVG. See it here: <http://jsdo.it/event/svggirl/>

SVG performance for this sort of complex animation used to be very lackluster, but hardware acceleration makes the world of difference. Hats off to the IE team for hardware-accelerated SVG. (You can read more about it here:

<http://blogs.msdn.com/b/ie/archive/2011/03/08/comparing-hardware-accelerated-svg-across-browsers-with-santa-s-workshop.aspx>.)

## SVG Edit

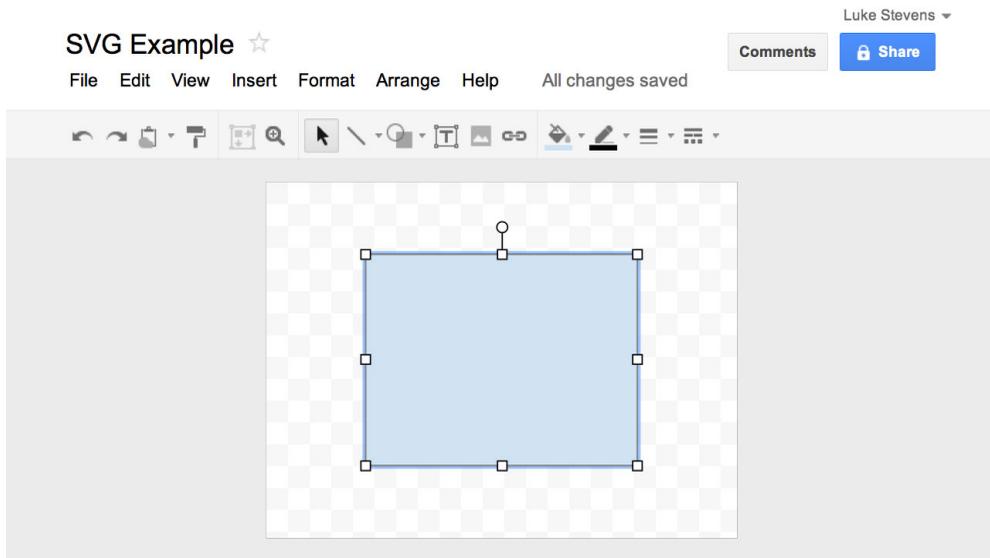


*Figure 11.3. SVG Edit is a SVG-powered drawing program that outputs... SVG.*

SVG Edit shows what can be done with client-side web technologies alone. It's an SVG- and JavaScript-based app for editing SVG. Download it here:

<http://code.google.com/p/svg-edit/> or try it live here: <http://svg-edit.googlecode.com/svn/branches/2.5.1/editor/svg-editor.html>.

## Google Docs



*Figure 11.4. SVG is used in some very high-profile situations, such as the Google Docs drawing program.*

Google Docs' drawing program uses SVG with a VML fallback. (Google also started using SVG for the graphs in Google Analytics in early 2012.)

## SVG Games

### SVG-oids

This re-creation of a classic 1979 arcade game shows what can be done with a little bit of SVG, JavaScript, and programming skill. Use the up and down keys to move, the left and right keys to rotate, and the spacebar to fire. Press the "n" key to play in a nebula.

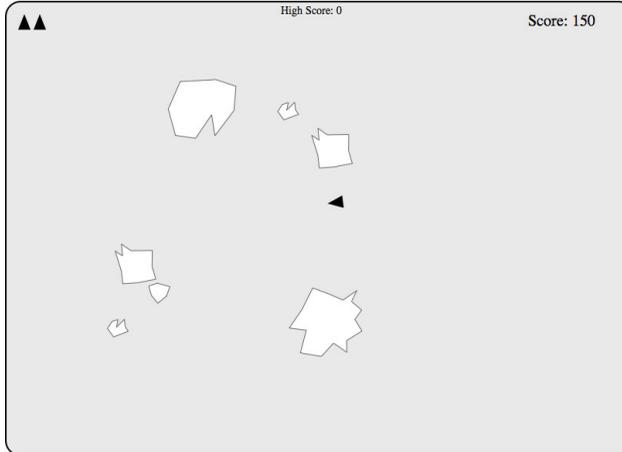


Figure 11.5. *SVG-oids* may be simple, but it demonstrates the interactive possibilities of SVG.

You can create games in SVG with JavaScript, but SVG for games hasn't caught on in the same way Canvas has. As part of their IE9 tech demos, Microsoft released a couple of *very* simple retro game examples:

- Asteroids in SVG: <http://ie.microsoft.com/testdrive/Graphics/SVGoids/Default.xhtml>
- A simple Helicopter game: <http://ie.microsoft.com/testdrive/Performance/Helicopter/Default.xhtml> with a background that reminds me of the Atari 2600.

## D3.js



*Figure 11.6. It's worth exploring the impressive D3.js examples online, such as this "streamgraph".*

D3.js (<http://mbostock.github.com/d3/>) is a "small, free JavaScript library for manipulating documents based on data" that uses SVG for some extremely impressive data visualizations. Check out the examples for more: <http://mbostock.github.com/d3/ex/>. Also see the 150+ annotated slides from a D3 workshop by D3 creator, Mike Bostock: <http://bost.ocks.org/mike/d3/workshop/>. It starts with a simple introduction of D3 and SVG and finishes with some impressive examples.

## Charts With Highcharts

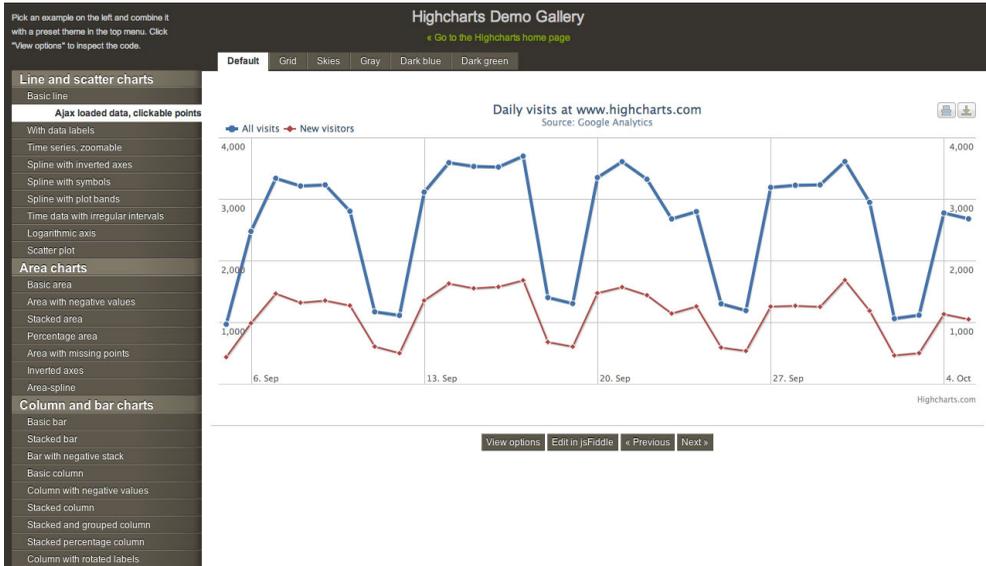


Figure 11.7. Highcharts has a lot of fans, and its flexible, well documented API makes it easy to use.

For a fully-featured SVG charting library, it's hard to go past Highcharts (<http://www.highcharts.com/>)—an SVG (and VML for legacy IE support) JavaScript-driven chart library. (Here they explain why they used SVG: <http://www.highcharts.com/component/content/article/2-news/12-highcharts-goes-svg>. Highcharts also gets a lot of love from developers too: <http://news.ycombinator.com/item?id=1847569>.)

## Raphael.js-powered Demos

Much of the current, real-world work with SVG is done with Raphaël which, as we saw earlier, offers a simple, cross-browser way to generate basic SVG.

## thirteen23

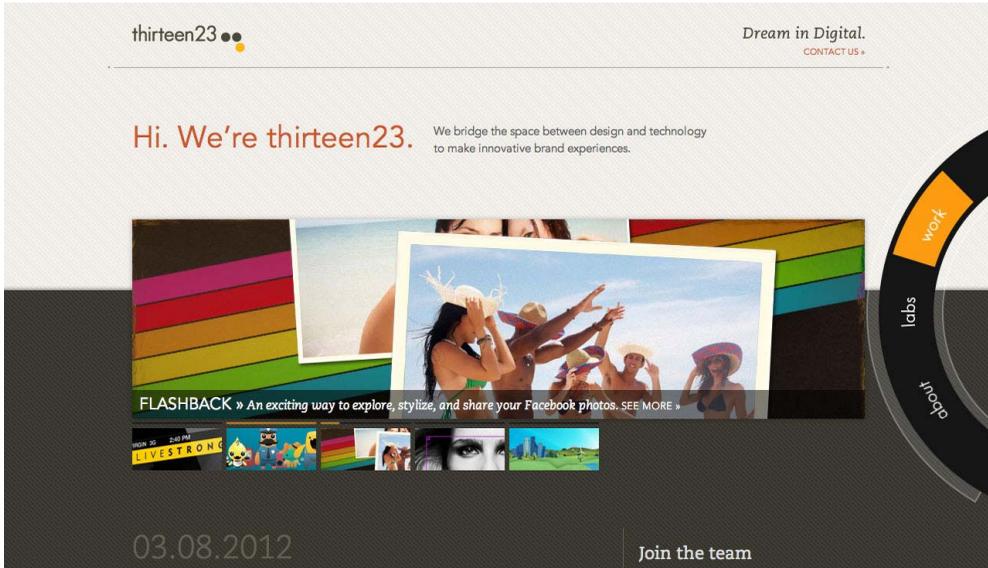


Figure 11.8. The animated, circular navigation for thirteen23 shows what a nice sprinkling of SVG can do.

thirteen23 (<http://thirteen23.com>) is a design consultancy in Austin, Texas that has done a great job using modern web technologies to pull off an impressive studio site. Their curved, smoothly-animated navigation is built using Raphaël. Click around to see the navigation in action (and watch the backgrounds change). Also note the lack of full page refreshes despite URL changes, and without using the `/#/` pattern. That's the HTML5 History API in action, which we'll touch on in the next chapter.

## Nissan Leaf



Figure 11.9. Once upon a time the Nissan Leaf site may have been a Flash affair, but it's all native web technology here.

The Nissan Leaf site (<http://www.nissanusa.com/leaf-electric-car/>) is a great example of what a smattering of SVG (using Raphaël) and a lot of modern JavaScript can do. The interface or style may not be to everyone's taste, but the point is the technology and the execution. We can do this kind of Flash-style interactivity now—no plugins required.

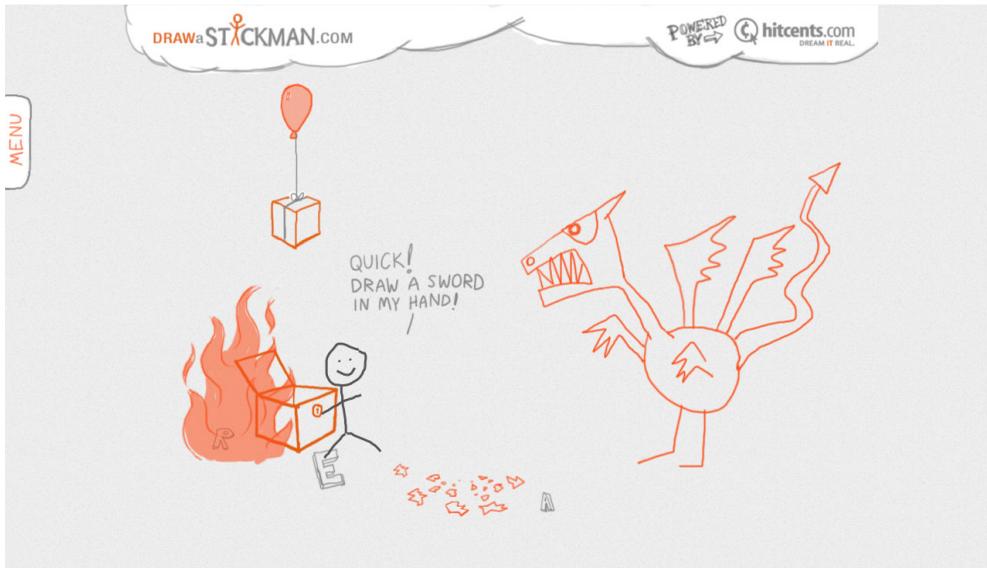
## Markup.io



Figure 11.10. Scribble all over your pages with SVG thanks to Markup.io

Markup.io (<http://markup.io/>) lets you draw vector lines on (and annotate) any web page with a simple bookmarklet. You can also publish and share your annotated pages. The drawing tools are SVG powered by Raphaël.

## DrawASTickman.com



*Figure 11.11. My stickman remained unflappably cheery in the face of a fire-breathing dragon.*

The <http://DrawASTickman.com> agency promo for Hitcents is, in the words of the creator, an “interactive website where visitors draw a stickman and take part in his animated adventure” which “became a viral success overnight, attracting millions of visitors from every part of the globe and winning numerous awards” (<http://www.hitcents.com/blog/post/making-drawastickmancom-part-1-birth-idea>). This brilliantly executed idea has had over 20 million visits, and uses Raphaël for the graphics.

## Election Results

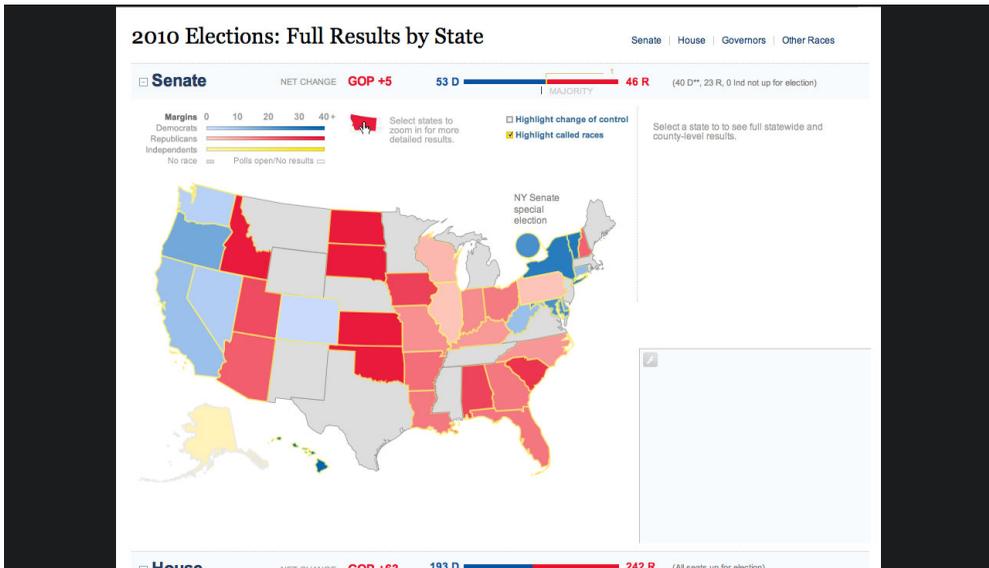


Figure 11.12. SVG is particularly useful for static, interactive vector graphics such as maps.

The Wall Street Journal used Raphaël in their 2012 US election map:  
<http://newsapps.wsj.com/elections2010/>.

## Visualizations

### It's All Connected: An Overview of the Euro Crisis

European leaders are dealing with growing debt problems that are rattling investors worldwide. Here is a visual guide to the crisis.

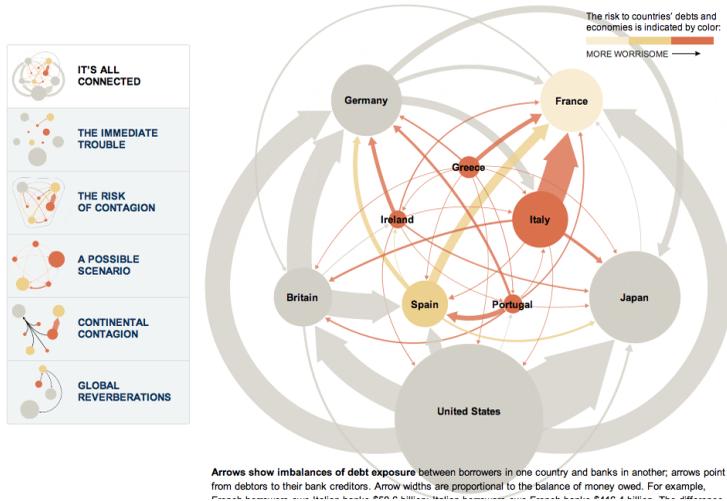


Figure 11.13. The New York Times interactive visualization is another great example of what SVG (and Raphaël) excel at.

The New York Times produced an excellent interactive graphic of the 2011 European debt crisis using SVG and Raphaël. SVG is excellent for these sorts of visualizations, and having major corporations such as the NYT and the WSJ using it is quite an endorsement.

## Working With SVG

Given cutting-edge browsers now supporting SVG files in `<img>` tags and even as CSS backgrounds, soon we'll get to use SVG files for background gradients, tab backgrounds (see this demo: <http://helephant.com/2009/08/12/svg-images-as-css-backgrounds/>), and other image elements where a single file can be reused and scaled as needed. Combine SVG's flexibility with CSS3's multiple backgrounds, and there could be some interesting possibilities on the horizon. For example, SVG would be great for styling the controls of a media player for `<video>` or `<audio>` elements.

This isn't just theoretical—designers are starting to consider SVG seriously, and provide tutorials to get us up to speed. For example, in “A Farewell to CSS3 Gradients”, Alex Walker looks at the patchy support for CSS3 gradients; suggests we consider SVG as an alternative; and provides a handy tutorial to do just that: <http://designfestival.com/a-farewell-to-css3-gradients/>.

## Responsive Web Design And SVG

Vector graphics may also prove quite helpful in responsive web design situations where we want to display crisp, lightweight interface elements on everything from the mobile to the desktop (and especially on super high resolution iOS and Android (3.0+) screens). Again this isn't theoretical; designers are getting their hands dirty with this right now. In Smashing Magazine's January 2012 article “Resolution Independence With SVG”, David Bushell looks at using SVG for interface elements: <http://coding.smashingmagazine.com/2012/01/16/resolution-independence-with-svg/>.

Vector UI elements are not a free responsive lunch, however. While it's tempting to think they can scale up and down effortlessly (they're vectors!), this isn't necessarily true. Large artwork scaled down to very small sizes can become a blurry mess, and small artwork scaled up can look spartan and devoid of detail, as this lengthy article on icons and SVG demonstrates: <http://www.pushing-pixels.org/2011/11/04/about-those-vector-icons.html>.

(Of course, CSS3's support for gradients, rounded corners, transformations and animation may place the proverbial pillow on SVG's face just as it comes out of life support for the third time. In fact, SVG filter effects are being ported to CSS and are already arriving in WebKit, as this article explains:

<http://updates.html5rocks.com/2011/12/CSS-Filter-Effects-Landing-in-WebKit>. Plus, check out these crazy demos of SVG-inspired CSS shaders: <http://www.adobe.com/devnet/html5/articles/css-shaders.html>.)

Despite the fact SVG has been kicking around for more than a decade, the web design community hasn't really given it a thorough workout to see what's possible. So with the renewed interest in all things web standards, the decline of Flash, the rise of responsive web design, and a reasonable baseline of browser support, maybe it's time to get experimenting again.

## SVG Gotchas

There are two key problems though. The first is performance: complex SVG is slow. Browser makers generally haven't paid much attention to SVG performance, as it was always the red-headed stepchild of web standards. Things are starting to change, however. For one, hardware acceleration helps dramatically, as Microsoft has demonstrated with IE9 and IE10. (Yes, *Microsoft* has not only caught up, but is now leading the way in some areas of web standards implementation. You're telling me.)

The other problem is tools. No one wants to sit around writing SVG markup by hand. There are drawing tools available however, such as:

- **Inkscape**, an open source, cross-platform vector drawing program that uses SVG natively: <http://inkscape.org>.
- **Adobe Illustrator** supports SVG, and you can read more about saving files as SVG here: <http://quintaldesigns.com/articles/svg-files-in-adobe-illustrator>. (Adobe also offers a HTML5 pack for Illustrator CS5, which extends its SVG support: [http://labs.adobe.com/technologies/illustrator\\_html5/](http://labs.adobe.com/technologies/illustrator_html5/). It even allows you to designate some elements be rasterized as a Canvas element, as described here: <http://rwillustrator.blogspot.com.au/2010/09/web-designers-rejoice-adobe-releases.html>. There's also a paid SVG kit plugin for **Adobe Creative Suite**: <http://svg.scand.com>.)
- **SVG-edit**, the "fast, web-based, Javascript-driven SVG editor" that's free and open source, which we looked at earlier. You can try it live here: <http://svg-edit.googlecode.com/svn/branches/2.5.1/editor/svg-editor.html>, or download it from <http://code.google.com/p/svg-edit/> (see the links to browser add-ons, too).
- And a variety of other tools listed here: [http://en.wikipedia.org/wiki/Scalable\\_Vector\\_Graphics#Software\\_and\\_support\\_in\\_applications](http://en.wikipedia.org/wiki/Scalable_Vector_Graphics#Software_and_support_in_applications).

You can also use JavaScript libraries such as Raphaël or jQuery SVG to draw SVG with JavaScript. Nevertheless, most vector drawing and animation for the web is done in the Flash IDE. What if Flash could export SVG?

## Flash Breathes Life Into SVG?

Interestingly, Adobe has recently released a Flash-to-HTML5 conversion tool that takes basic animation from a Flash FLA file and converts it to SVG, CSS3

and JavaScript. It's an experimental FLA-to-HTML tool called Wallaby which relies heavily on SVG and WebKit-only CSS3: <http://labs.adobe.com/technologies/wallaby/>.

Of course, that's not *really* HTML5, but we'll give them a buzzword-pass on this one. Here's Adobe's John Nack (<http://blogs.adobe.com/jnack/2011/03/wallaby-flash-to-html5-conversion-tool-now-available.html>):

*Adobe's job is to help you solve problems, not to get hung up on one technology vs. another.*

*Millions of people have honed their Web animation skills in Flash, and now their customers want content that can run anywhere, including on non-Flash-enabled devices. Accordingly Adobe's releasing "Wallaby," an experimental Flash-to-HTML5 conversion tool.*

Flash export to SVG may be another case where Flash, of all things, drives the adoption of new web standards. It makes sense, and the sentiment has been around for a while (Jonathan Snook expressed hopes for something similar in 2009: <http://snook.ca/archives/opinion/adobe-html5-canvas>). As John Nack says, Flash is where millions of people do web animation and vector graphics. And this is one way for designers to get their (basic) Flash animations to work on iOS, and soon Metro IE10 and all future Android 4.x devices.

But this assumes the translation from Flash to basic SVG is acceptable in terms of both similarity and performance. Currently this is just experimental technology from Adobe, so we shouldn't expect too much. But if nothing else we may eventually see banner ads animated with SVG (or Canvas) targeted at iOS and Android devices when performance becomes acceptable. *Oh joy.*

The question isn't whether SVG can replace Flash, but whether it can carve out a niche of its own. Raphaël has certainly helped it carve out that niche, and in the next few years we'll be able to do more with SVG natively.

Can SVG generate enough interest in the design community after crying wolf in the early '00s? Ironically, it may not be Flash that distracts the web community from SVG, but other web standards. With so much activity around web standards development these days, it's got a fight for attention on its hands. Is

SVG destined to always be the bridesmaid, and never the bride? Or will the guys who wrote the huge tomes on SVG all those years ago get to dust off their work for a belated second edition? Time will tell.

In the meantime, it's easy to get started with Raphaël, so why not give it a try? Check out <http://raphaeljs.com>.

# THE TRUTH ABOUT HTML5 WEB APPS, MOBILE & WHAT COMES NEXT

This book is primarily about HTML5 for web designers. And what HTML5 delivers for us is... well, mixed.

But it's not fair for that impression to characterize the entire HTML5 spec. Many of the features now in HTML5 began life in the Web Applications 1.0 spec (<http://www.whatwg.org/specs/web-apps/2005-09-01/>), and it shows. HTML5 wasn't really written for designers. Instead it was written more for developers. And from a development point of view, HTML5 looks far more interesting, even for us as designers, as we'll see in a moment.

So, in this chapter, we're going to take a quick run through some of the important web app-oriented features of HTML5 which are, in some ways, the real meat of HTML5.

First though, let's take a quick look at the (rapidly changing) browser landscape for HTML5 web app development.

### HTML5 Web App Browser Support

For all the hype, HTML5 browser support for web app features (on the desktop at least), is still a bit of a downer. But we can't blame that on the spec. Ian Hickson and the WHATWG have bent over backwards to make HTML5 as implementor-friendly as possible, documenting browser behavior in a way that's never been done before.

Legacy IE versions (particularly IE8—the last version of IE for XP) will still be with us for years to come. (And IE users aren't exactly rushing to upgrade—see Ars Technica's analysis of browser upgrade patterns here: <http://arstechnica.com/web/news/2011/06/may-browser-market-share-microsoft-and-mozillas-continuing-chrome-conundrum.ars>.) IE8 is rapidly becoming the (now thankfully dead) IE6—that hard to shake, pain in our collective behinds, which will hang on as stubbornly as Windows XP does. (Here are the worldwide operating system trends, but always check your own analytics data to see what's relevant for your audience: <http://gs.statcounter.com/#os-ww-monthly-201102-201202>.)

But the next decade is actually looking much better than the previous one. Chrome and Firefox are now on release cycles measured in *weeks*. As Jeff Atwood said in his post “The Infinite Version” (<http://www.codinghorror.com/blog/2011/05/the-infinite-version.html>):

*Chrome is so fluid that it has transcended software versioning altogether.*

Microsoft, desperate not to be left behind, is not only catching up as fast as they can, they're innovating too—from hardware accelerated Canvas and SVG, to a proper CSS layout system in IE10. (IE10 will be the first browser to have CSS Grid Layout—check out the spec here: <http://dev.w3.org/csswg/css3-grid-align/>.) Microsoft is also speeding up their browser releases, aiming for annual version updates (see: [http://www.computerworld.com/s/article/9215766/Microsoft\\_quickens\\_browser\\_pace\\_with\\_IE10\\_goes\\_for\\_annual\\_upgrades](http://www.computerworld.com/s/article/9215766/Microsoft_quickens_browser_pace_with_IE10_goes_for_annual_upgrades)), and pushing HTML5 hard for Metro style apps in Windows 8. They're also pushing automatic upgrades to get Windows users on the latest browser for their operating system (as Ars Technica discusses here: <http://arstechnica.com/microsoft/news/2011/12/microsofts-new-automatic-update-plan-could-finally-spell-the-end-of-ie6.ars>).

In the meantime, HTML5 polyfills will keep being refined and enhanced to fill the gaps as best we can until legacy IE (and WinXP) dies out: <https://github.com/Modernizr/Modernizr/wiki/HTML5-Cross-browser-Polyfills>.

Browser updates and innovation are as important as the spec itself. As Ian Hickson says, a spec that doesn't get implemented exists only as fiction. But this

raises other issues: how do you develop for “modern” browsers when they’re being updated every couple of months? Part of the answer is feature detection, which we’ll look at shortly.

Then there’s mobile...

## HTML5 On Mobile: WebKit And Beyond

If HTML5 web app development in the here-and-now has one saving grace, it’s mobile (specifically iOS and Android) web app development. The support for HTML5 features in WebKit (the saucy minx behind iOS’ and Android’s browsers) is solid, and improving all the time. But we can’t assume the mobile web means WebKit even if it dominates, just as we didn’t develop exclusively for IE on the desktop even when it dominates.

In fact, we can’t even assume “WebKit” refers to a consistent platform, either. In A List Apart’s “Smartphone Browser Landscape” article from December 2010, by Peter-Paul Koch writes (<http://www.alistapart.com/articles/smartphone-browser-landscape/>, emphasis in the original):

*There is no WebKit on mobile. I tested nine mobile WebKit-based browsers and they all behave differently. Not wildly so: Baseline CSS support is good, and JavaScript is definitely workable. Still, each one has its problems and strong points.*

*Because of this variability, it’s important to test your websites in as many WebKit-based browsers as you can. Don’t assume your website will work on the Android or BlackBerry WebKit-based browsers just because it works in Safari.*

That was as true then as it is now, if not more so. For example, there are two different Android browsers from Google for Android 4.0—the stock Browser app, and the new Chrome for Android (see: <http://googleblog.blogspot.com.au/2012/02/introducing-chrome-for-android.html>), and that’s not taking into consideration Android OS fragmentation, which is in and of itself a bit of a nightmare.

Even talking about “mobile” as though it’s synonymous for iOS and Android is problematic. The mobile world is so vast—including the “smartphone” market, which has plenty of cheap, under-powered Android devices and legacy iOS devices—it’s a mistake to generalize at all. There’s also Opera Mobile and Mini, and Firefox for Mobile. (In fact, there’s a whole Mozilla mobile platform coming, which we’ll discuss below.) And then there’s Internet Explorer Mobile.

## Mobile Is A Moving Target: Microsoft’s Big Push

While the current mobile browser landscape is complicated, to say the least, it’s also in a state of flux. Big players—well, one big player specially—are only just making a serious push into the market, and that’s Microsoft, with their Windows Phone partnership with Nokia.

Microsoft’s Windows Phone 7 platform launched in late 2010, and shipped with our old friend—and by friend I mean sworn enemy—IE7 (well, IE7 and a half). Microsoft thankfully started rolling out IE9 Mobile in late 2011 as part of their Windows Phone 7.5 (a.k.a. “Mango”) update.

IE9 Mobile is more or less the same as IE9 on the desktop, and it has a lot of good things going for it—hardware accelerated Canvas, for one (for more, see the list here, including CSS3 and SVG support: [http://windowsteamblog.com/windows\\_phone/b/wpdev/archive/2011/09/22/ie9-mobile-developer-overview.aspx](http://windowsteamblog.com/windows_phone/b/wpdev/archive/2011/09/22/ie9-mobile-developer-overview.aspx)).

Microsoft has unified its desktop and mobile browser codebase, and given they are pushing HTML5 hard across the board (especially in IE10), *and* can push out mobile browser upgrades quite efficiently (unlike the desktop), we’ll hopefully see IE Mobile on par with WebKit in terms of HTML5 support in the not-too-distant future.

In the meantime, however, IE9 Mobile is another example demonstrating that the modern mobile world isn’t just iOS and Android. It also introduces a new set of assumptions for mobile web apps about what can and can’t be supported. For example, there’s no HTML5 Application Cache (i.e. offline mode). The list of what’s missing in IE9 Mobile is much the same as what’s missing in IE9 on the desktop (see: <http://people.mozilla.com/~prouget/ie9/>).

We also need to keep mobile use in perspective. Mobile, and particularly responsive web design may be, *like, so hot right now*, but if only 8% of your audience uses smartphones to access your site, you need to weigh up how you spend your time for that 8% versus the other 92% of users. (On the other hand, if your numbers are the reverse, go nuts!) Check your stats in Google Analytics, or (shameless plug ahoy!) Ninja For Google Analytics at <http://itsninja.com>, which I designed for web designers and developers so you can have the relevant Google Analytics data at your fingertips.

## Boot To Gecko: Mozilla's Ambitious Mobile Platform, And The WebAPI

While companies like Microsoft try and push their own platform, a part of which involves the web, Mozilla has been working hard at turning the web itself into a mobile platform, especially with its Boot to Gecko (B2G) project. In their own words (<http://hacks.mozilla.org/2012/02/mozillas-boot-to-gecko-the-web-is-the-platform/>):

*Mozilla's Boot to Gecko (B2G) is about building a complete, standalone operating system for the open web. It aims at making web technologies the number one choice for applications on desktop and mobile, and we believe it can displace proprietary, single-vendor stacks for application development. And we have made some exciting progress that we want to share with you!*

This ambitious project aims to deliver “HTML5” (I use the term loosely) devices—particularly phones:

*Enabling HTML5 devices running on the Open Web which can deliver smartphone capabilities at feature phone prices.*

This is “HTML5” the marketing buzzword; as a synonym for all things “Open Web”. It's actually a post-HTML5 project that uses some of the functionality detailed in the HTML5 spec, and a whole lot more Mozilla has been developing under their WebAPI banner (see <http://hacks.mozilla.org/2011/08/introducing-webapi/>)

and <http://hacks.mozilla.org/2012/02/mozilla-and-the-mobile-web-api-evolution/> for more).

The WebAPI project has set its sights high, and aims to fill the many large gaps in using web technology as a mobile platform, and includes APIs for things like a phone dialer, SMS functionality, an address book, camera access, device settings, games, and more.

This means you can have a phone powered entirely by web apps, and you can see how they run simply by *viewing the source*, just like any other web app or web site. Mozilla have been offering hands-on demos of just such a phone. For example, The Verge reported from the Mobile World Congress in February 2012 (<http://www.theverge.com/2012/2/27/2827659/mozillas-boot-to-gecko-project-the-internet-is-your-phone-hands-on>) that:

*Gecko is, in essence, a phone operating system that is entirely web- and HTML5-based. From the moment you turn the phone on, everything you see is HTML5. Even the dialer uses Mozilla's "telephony APIs," and is itself web-based. There are no native apps, just a series of the most impressive bookmarks you've ever seen. [...]*

*Sending messages, taking pictures, playing Cut the Rope, browsing the web, and nearly everything else we tried worked correctly, if not always gracefully. It was actually really hard to believe that we were using an entirely web-based device—we kept asking if they were lying, and it wasn't really HTML5. Of course, there was an easy way to prove it: you can see the source code of any app at any time, to see exactly what's behind what you're seeing.*

Again, this is “HTML5” in the broadest sense of existing and new (particularly Mozilla-invented) web technology, not the HTML5 specification per-se. Nevertheless, it demonstrates the incredible potential for web technology beyond HTML5, and the power of the web as a platform in its own right.

It also shows development of web technology is not purely a matter of the WHATWG vs the W3C, as we saw way back in chapter one. It's possible for players such as Mozilla to do plenty of innovation in their own time (and on their

own dime), and (as Mozilla is doing) take their work to the W3C to ensure it becomes a patent-free standard all can use and implement.

## HTML5 Mobile Compatibility

For more on the current state of HTML5 mobile compatibility, Maximiliano Firtman has put together an excellent chart with 15 mobile browsers, and a variety of HTML5 and related web technologies and APIs, that's available here:

<http://mobilehtml5.org/>.

That wraps up our quick journey through HTML5 and mobile, and now let's get back to HTML5 and web app development.

## HTML5-Powered Content Management

One key reason designers should have at least some interest in HTML5's web app features on desktop and mobile is Content Management Systems. CMSs are the one category of web apps (broadly speaking) we *all* rely on for our client or company work.

For example, it would be great to see CMSs take advantage of the History API for fast page loads, local storage (for auto-saving form entries, perhaps), and maybe offline functionality for mobile blogging/content editing. (We'll touch on these features below.) There's also the Mozilla HTML5 Image Uploader (<http://hacks.mozilla.org/2010/02/an-html5-offline-image-editor-and-uploader-application/>) that uses a bunch of HTML5 technology such as the offline application cache, local storage, Canvas and drag and drop, and serves as a good indicator of how HTML5 can enhance the CMSs we (and our clients) use day in, day out. The sooner we suggest these features, the sooner they'll be implemented.

## The JavaScript Age

While we may spend a lot of time in the trenches wrangling CMSs, it's worth keeping an eye on the big picture too. The overarching theme of these new features—and possibly the future direction of the web itself—is they're all about JavaScript APIs, and not HyperText, per-se. In a sense, HyperText is now the tail wagged by the JavaScript dog, especially in terms of web applications.

The problem of marking up documents has been solved. Over the next decade or so we'll see major improvements in writing applications, as HTML5 (and related) features for web apps become baseline standards. After all, HTML5 is largely a web apps spec from the mid 00s.

This is a (small) part of the reason we're perhaps entering the JavaScript Age, as Mike Driscoll puts it in “Node.js and the JavaScript Age”

(<http://metamarketsgroup.com/blog/node-js-and-the-javascript-age/>):

*[We need to shift] our view of the server from a document courier (HTML Age), or a template renderer (LAMP Age), to a function and data shipper. The principal role of the server is to ship an application to the client (JavaScript), along with data (JSON), and let the client weave those into a DOM. [...]*

*The JavaScript age brings us closer to a web that is not a global digital library, but a global digital nervous system, whose implications we are only beginning to apprehend.*

Or, put another way, by Jeff Atwood in “The Principle of Least Power”

(<http://www.codinghorror.com/blog/2007/07/the-principle-of-least-power.html>):

*Atwood's Law: any application that **can** be written in JavaScript, **will** eventually be written in JavaScript.*

The question is, what will this mean for the humble web page?

## JavaScript Killed The HTML Star

In a sense, we have been exploring this new frontier since the whole “Web 2.0” thing kicked off. So we shouldn’t be surprised that the thrust of new developments are JavaScript-related.

The humble web page has been incredibly robust and resilient. With billions and billions of them out there, and billions more being made, they’re not going away any time soon. But the hugely inefficient *click - full page refresh - click - full page refresh* paradigm in high profile sites may not be around for long in light of things such as the History API, which we’ll look at in a moment.

We’ve achieved some pretty incredible things on top of the basic web page. Uber-nerds get excited about complex functionality (as is happening now with web apps-oriented HTML5 functionality), which then gets abstracted into a library, plugin, or framework and becomes almost trivial for designers to drop in to their sites.

While this is wonderful for us, it does mean the lines between web *page* and web *application* will continue to blur. Is a modern web site a JavaScript/AJAX/HTML5-powered *application* that serves content, or a collection of simple, linked pages, or something in between? Right now the answer is “all of the above”, but it will be interesting to see at what point the amount of JavaScript on our site makes it more “app” than “page”.

For example, consider the amount of functionality that can be bolted onto a traditional web page through JavaScript. JavaScript for animation; JavaScript for SVG (with Raphaël, as we saw in chapter eleven); JavaScript for page state (which we’ll look at below); JavaScript for A/B testing your design; JavaScript for new CSS layout engines (e.g. <http://code.google.com/p/css-template-layout/>); JavaScript for CSS preprocessing; JavaScript for graphics and games (Canvas and WebGL, as we saw in chapter nine); JavaScript for audio and video controls (chapter ten); and even JavaScript for decoding MP3s (<http://hacks.mozilla.org/2011/06/jsmad-a-javascript-mp3-decoder/>, which works around the patent issues we discussed in chapter nine, and is pretty incredible in its own right).

Phew.

JavaScript is the wrapper around (or enabler of) all this functionality, some of which is HTML5, and some of which is not. But how long until we go all the way—if indeed we should—and launch bare-bones HTML pages with everything else abstracted into essentially one big JavaScript app? If some web apps are now being delivered as simply JSON data and a client-side JavaScript application (as we touched on above), why not web pages?

Well, there's SEO, but the situation is rapidly changing there. In November 2011 Google's Matt Cutts said the Googlebot was already spidering JavaScript/AJAX comment systems, at the very least (<http://searchengineland.com/google-can-now-execute-ajax-javascript-for-indexing-99518>); Google can crawl the (problematic) hashbang (#!) URL format; and the HTML5 History API provides new possibilities we'll look at below. (Apart from SEO, issues like maintenance are, however, quite legitimate objections!)

It's clear which way the wind is blowing. Think how far we've already come from the days when JavaScript was synonymous with horrendous DHTML and clunky roll-over scripts. Now it has become the de facto programming language of the web.

Nevertheless, we shouldn't get *too* distracted by the technology. It's fun to gaze into the crystal ball, but the evergreens such as great copy, a great user experience (especially when measured in hard numbers like conversion and/or engagement rates), and designing to generally *get the hell out of the way* will still matter above all else. Some things just never go out of fashion. (We'll touch on this a little more when we look at Performance Based Design in the final chapter.)

But before we get there, let's quickly run through these new web app-oriented features and how we can detect them (along with some resources for further reading.)

## Modernizr, When Can I Use..., And Polyfills

Browser releases have become incredibly fast. As we touched on earlier, Chrome and Firefox aim to roll out updates in *weeks*. This alone makes it incredibly

difficult, if not outright foolhardy, to try and detect which browser version gets which functionality.

Browser detection has never been a great idea. You still occasionally run into antiquated websites that tell you to “upgrade” your bleeding-edge version of Chrome or Firefox to a “modern” browser like Internet Explorer 7, because that’s all the developers could foresee at the time. It’s hard to detect browsers that don’t yet exist.

## Modernizr

Instead, *feature* detection is now all the rage, and that’s what scripts like Modernizr do. They don’t *add* any functionality. They simply tell you what features are supported in a given browser so we can tailor our page as appropriate.

Modernizr does this in one of two ways:

1. By adding a class name to the `<html>` element (especially useful for CSS3 features) so we can write css for `.coolfeature {}` or fallback styles for `.no-coolfeature {}`.
2. Through a global JavaScript object that contains properties for each HTML5 (and related) feature. Where those properties evaluate as true or false reflects whether the feature is supported or not. The built-in YepNope.js library allows conditional loading of supported scripts (yep, the feature is supported) and polyfills (nope, it’s not, so load this polyfill).

Check it out at: <http://modernizr.com>.

There’s a full tutorial on [HTML5doctor.com](http://html5doctor.com/using-modernizr-to-detect-html5-features-and-provide-fallbacks/): <http://html5doctor.com/using-modernizr-to-detect-html5-features-and-provide-fallbacks/>.

## When Can I Use...

This is great for detecting features on a browser-by-browser basis, but how can we know what support is like for a given feature in the first place? To get an idea of global browser support for a given feature, I recommend Alexis Deveria’s incredibly useful <http://caniuse.com> (which I’ve referred to in previous chapters).

## Polyfills

Fallback functionality for unsupported browsers can be *sort of* enabled with “polyfill” scripts and hacks. Here’s an excellent, near-exhaustive list of what’s available: <https://github.com/Modernizr/Modernizr/wiki/HTML5-Cross-browser-Polyfills>. Keep in mind these are rarely get-functionality-for-free cards; there’s always compatibility and performance issues to consider.

## HTML5 Web App APIs

Okay, let’s jump into the new HTML5 (and thereabouts) web app functionality.

(Browser stats in this chapter come from <http://caniuse.com>, which only goes back to IE 5.5, Firefox 2, Chrome 4, Safari 3.1, Opera 9, iOS 3.2, Opera Mini 5, Opera Mobile 10, and Android 2.1. Windows Phone 7’s current browser is IE9 with slight differences. I’m leaving out Opera’s mobile browsers and Firefox Mobile for simplicity’s sake, but keep in mind the rapid changes in the very broad mobile market we discussed earlier.)

## History API (pushState)

Let’s start with the HTML5 History API (also referred to as “pushState”). URLs have been abused to all hell for fun and profit in the AJAX era. This is especially true with the hash-bang (!) approach you may have seen on Twitter, Facebook, and sites such as Gawker. (Typing <http://twitter.com/lukestevens> actually returns <http://twitter.com/#!/lukestevens>, but Twitter is moving away from this behavior.)

Some people think this is, like, *the worst thing ever*, while others think it’s the price of progress. See the debate, with links, here: <http://danwebb.net/2011/5/28/it-is-about-the-hashbangs>. (That’s from Dan Webb, who’s in charge of *undoing* Twitter’s hashbang URLs, as he tweeted here: <https://twitter.com/#!/danwrong/status/171680703824662528>.)

In any case, the History API should go some way to solving this. With the History API we can still load in a chunk of new content into your page with AJAX (or similar), but we can also update the URL in the user’s location bar

(*and* browser history) to make the whole process look like a very fast page request.

Mind you, it takes some work to fake the whole backward/forward page navigation thing. You can read Mark Pilgrim’s detailed write up and tutorial here: <http://diveintohtml5.info/history.html>. SEOmoz also covered the History API in “Create Crawlable, Link-Friendly AJAX Websites Using pushState()” (<http://www.seomoz.org/blog/create-crawlable-link-friendly-ajax-websites-using-pushstate>).

This would be a handy addition to the modern CMSs we offer our clients (at the very least). We’d get fast AJAX page loads without confusing not-necessarily-technical clients, who’d still see the new URL and have their back button work predictably.

It’s also something we can consider implementing in a progressive way. As of writing, there’s no IE9 support (though it is in IE10); Chrome support is good; Safari support is buggy; and Opera support only arrived in 11.5+. (On mobile, iOS 4.2-4.3 support was buggy, iOS5 support is solid, and Android inexplicably dropped support in 4.0.)

But there is History.js (<https://github.com/balupton/History.js>), which aims to provide fallback support for unsupported browsers and smooth out the quirks in current implementations. The creator of History.js Benjamin Lupton discusses the pros and cons of the different approaches to handling URLs in “Intelligent State Handling” (<https://github.com/balupton/history.js/wiki/Intelligent-State-Handling>).

(Remember though, hashbang URLs are forever, even if only used as a fallback for older browsers. If someone uses a hashbang URL to link to your site, then that’s a URL that needs to be maintained indefinitely. Falling back to full page loads may be a better way to go.)

For current browser support stats for the History API, see:

<http://caniuse.com/#feat=history>.

## HTML5 Web Storage (And JavaScript-Rendered CSS)

Web Storage (also known as “localStorage”) is often described as “cookies on steroids”, given it can store up to 5MB of data (key/value pairs) on the client. Unlike cookies, the data isn’t automatically sent back to the server; and it doesn’t have an explicit expiry date. (Web Storage was originally part of the HTML5 spec, but has since been spun out into its own specification, still edited by Ian Hickson: <http://dev.w3.org/html5/webstorage/>.)

This kind of storage could be used to save web app data locally, whether it’s a saved game state (so the user can resume where they left of), or a document the user is working on.

In terms of web design, we could use localStorage to save the output of a CSS pre-processor. In the past few years CSS pre-processors (such as LESS and SASS) have become all the rage, offering advanced CSS syntax and features such as variables, mixins, and better inheritance. You write the code using the new syntax, and the software spits out normal CSS a browser can understand. You can output the CSS as a one-off, or automatically on the server.

You can also do it client-side with JavaScript and Less.js (<http://lesscss.org/>). The Less.js script uses Web Storage to cache the outputted CSS, making subsequent requests for the CSS extremely fast (as described somewhat enthusiastically here: <http://fadeyev.net/2010/06/19/lessjs-will-obsolete-css/>).

This is a pretty profound development. We can now use JavaScript to generate CSS however and whenever we like, and store it locally. No more faffing about with server-side scripts to do the parsing, during development at least. (For production it should be compiled to normal CSS, otherwise those without JavaScript won’t get your CSS either, and as we covered in chapter four, that’s A Bad Thing™.) It’s simple, powerful, and works right now.

(By the way, I’m not advocating LESS over any other flavor of CSS pre-processor. If you decide to go down this route, use whatever floats your boat.)

For more on Web Storage, see:

- Mark Pilgrim’s chapter from Dive Into HTML5: <http://diveintohtml5.info/storage.html>.

- Opera’s “Web Storage: easier, more powerful client-side data storage” article: <http://dev.opera.com/articles/view/web-storage/>.
- Mozilla Developer Network’s “DOM Storage” article: <https://developer.mozilla.org/en/dom/storage>.
- This Mozilla blog post, which looks at Web Storage for saving the contents of a form (to avoid data loss when the browser crashes, the tab is accidentally closed, and so on) and touches on fallbacks for older browsers: <http://hacks.mozilla.org/2011/04/using-client-side-storage-today/>.
- There’s also been recent debate about the pros and cons of localStorage, with Chris Heilmann from Mozilla writing “There is no simple solution for local storage” (<http://hacks.mozilla.org/2012/03/there-is-no-simple-solution-for-local-storage/>) and John Allsopp responding with “localStorage, perhaps not so harmful” (<http://www.webdirections.org/blog/localstorage-perhaps-not-so-harmful/>).

Web Storage is currently supported in IE8+ and all other modern desktop browsers (FF 3.5+ Safari 4+, Chrome 4+, Opera 10.5+), and mobile browsers (iOS 3.2+ and Android 2.1+). (For current usage statistics, see: <http://caniuse.com/#feat=namevalue-storage>.)

## Database Storage

Web Storage sounds great, for some kinds of data. What about database storage client-side?

Politics, that’s what.

The no-longer-maintained Web SQL Database spec (<http://www.w3.org/TR/webdatabase/>) has been implemented in some browsers using SQLite (Safari, Chrome and Opera on the desktop), along with Mobile Safari and Android. But Microsoft never implemented it, and Mozilla is philosophically opposed to it (see: <http://hacks.mozilla.org/2010/06/beyond-html5-database-apis-and-the-road-to-indexeddb/>).

Mozilla is now pressing ahead with an alternative known as IndexedDB in Firefox 4+. Microsoft has expressed interest in it (and will be supporting it in IE10), and Google has started implementing it in Chrome 11+. There’s no Safari or Opera support as of writing, and mobile support is non-existent.

It will be a long time before there's any widespread, universal client-side database storage on the web. But the no longer maintained Web SQL Database may be an option for WebKit-focused mobile web app development.

## HTML5 Offline (Application Cache)

HTML5 allows developers to keep their web apps (or sites) running even if the client is offline—a common concern in the mobile world where a lost connection is just a tunnel away.

How? By specifying which URLs the browser should (and shouldn't) cache in a *manifest* file, which you reference by using the `manifest` attribute on the `<html>` element for every page of your web app.

It's one of those simple-in-theory-but-complicated-in-reality features that puts it beyond the scope of this book. So if you want to know more, check out:

- The web developer edition of the HTML5 spec, which has a lengthy explanation and tutorial: <http://developers.whatwg.org/offline.html>.
- Mark Pilgrim's chapter in Dive Into HTML5, which covers this feature in considerable detail, including debugging information: <http://diveintohtml5.info/offline.html>.
- Dev.Opera's handy introduction: <http://dev.opera.com/articles/view/offline-applications-html5-appcache/>.
- Steve Sounders has a nice write up of how the Application Cache (or "AppCache") can be improved, given developers currently have some pretty serious issues with it: <http://www.stevesouders.com/blog/2011/10/03/improving-app-cache/>.
- Atul Varma from Mozilla has also written a thorough post on the problems of developing offline apps in "The Challenges of Developing Offline Web Apps" (<http://www.toolness.com/wp/2011/06/the-challenges-of-developing-offline-web-apps/>).

Bringing some of these features together (along with the Geolocation API that we'll look at soon) can create robust, mobile, HTML5-driven web applications. Now we just need to wait for the desktop to catch up.

And by “desktop” I mean “Internet Explorer”, which doesn’t support HTML5’s offline functionality in IE9 down, but does support it in IE10. The other desktop browsers do support it (FF 3.5+ Safari 4+, Chrome 4+, Opera 10.6+), as does (importantly!) iOS 3.2+ and Android 2.1+. (For current usage statistics, see: <http://caniuse.com/#feat=offline-apps>.)

## Geolocation API

Geolocation on the web isn’t new. Plenty of sites use your IP address to work out your location (at least at the country level) so they can:

- serve you region-specific ads
- lock you out of certain services (as anyone living outside the US knows all too well!).

It can even be done client-side when you load any Google AJAX API (which you may already use to load jQuery—see <http://googleajaxsearchapi.blogspot.com/2008/08/where-is-my-current-user.html>) or server-side using something like <http://www.ip2nation.com/>.

The good news is you don’t need permission to do old-school geolocation based on your users IP. The bad news is it’s not always that accurate.

The new Geolocation API, on the other hand, tries to use whatever location data is available, including:

- GPS
- wireless network (as recorded by Google Street View)
- cell tower locations
- IP address (it’s agnostic as to where the data comes from).

It can then provide details about latitude, longitude, altitude, direction, speed, even accuracy data (when it’s available). Previous locations can even be cached (to map out a trip, for example).

But you can forget about getting this data without anyone knowing. The device *must* ask for your permission before it can use it.

The Geolocation API isn't part of HTML5 per se. (You can see the spec here: <http://dev.w3.org/geo/api/spec-source.html>.) But it's very cool, and opens up some very interesting possibilities with mobile web sites. The spec suggests:

- blogging (or status updates) with location data attached
- showing a user's position on a map
- turn by turn navigation
- tour guide web apps
- and location-specific weather or news widgets.

And remember: it all happens in the browser.

The Geolocation API is relatively well supported in all modern desktop browsers (IE9+, Safari 5+, Firefox 3.5+, Chrome 5+, Opera 10.6+), as well as iOS 3.2+ and Android 2.1+. (For current usage statistics, see: <http://caniuse.com/#feat=geolocation>.)

For more information, see:

- Mark Pilgrim's Dive Into HTML5 Geolocation chapter: <http://diveintohtml5.info/geolocation.html>.
- The "A Simple Trip Meter using the Geolocation API" tutorial [http://www.html5rocks.com/tutorials/geolocation/trip\\_meter/](http://www.html5rocks.com/tutorials/geolocation/trip_meter/).
- Firefox's user guide to "Location-aware browsing" <http://www.mozilla.com/en-US/firefox/geolocation/>.

## Other APIs That Are Completely Beyond Me But May Be Of Interest To You

Other new APIs in and around HTML5 allow developers to, in the words of Remy Sharp (Introducing HTML5, 2011, p216):

*You can now create multi-threaded, multi-window, cross-domain, low-latency, real-time thingymegiggies using the simplest of string-based communication methods. Now go build something awesome.*

These APIs include **Cross-document messaging** (IE8+ up, all modern desktop and WebKit mobile browsers), which allows documents on different domains to

send information back and forth. This could be useful for a widget on one site that needs to pull data from another domain. This (somewhat dated) MSDN article from 2008 provides a useful overview of the issues and technology:

[http://msdn.microsoft.com/en-us/library/cc511311\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/cc511311(v=vs.85).aspx).

**Web Workers** (IE10+, Firefox 4+, Safari 4+, all recent Chrome and Opera, limited mobile) offers an API to run scripts concurrently in the background, instead of the one-at-a-time-while-everything-freezes-until-its-done approach we have now. Wikipedia has a useful, brief overview: [http://en.wikipedia.org/wiki/Web\\_Workers](http://en.wikipedia.org/wiki/Web_Workers).

**Web Sockets** (IE10+, Firefox 6+, Chrome 14+; mixed support elsewhere primarily due to security issues) allows efficient bi-directional communication between browser and server. At its most basic, this could be handy for a chat room application. For more resources, see: <http://stackoverflow.com/questions/4262543/what-are-good-resources-for-learning-html-5-websockets>.

The **File API** (experimental or partial support in IE10+, Firefox 3.6+, Safari 5.1+, Chrome 6+, Opera 11.1+, no iOS, Android 3.0+) lets us read and write files and directories, as demonstrated in this neat music player:

<http://antimatter15.github.com/player/player.html>. There's also a useful tutorial here: <http://www.html5rocks.com/tutorials/file/filesystem/>.

Oh, and there's also the **Drag and Drop API** (IE5+, Firefox 3.5+, Chrome 4+, Safari 3.1+, Opera 12+, no mobile) which is, according to *every single person who has ever come into contact with it* (including Ian Hickson, who added it to the spec), pretty horrible. (Seriously, some people truly despise it:

[http://www.quirksmode.org/blog/archives/2009/09/the\\_html5\\_drag.html](http://www.quirksmode.org/blog/archives/2009/09/the_html5_drag.html).)

Drag and Drop (DnD) was added to IE5—yep, 5—by Microsoft in 1999, and other browser vendors also implemented it. So Hickson reverse-engineered it (ala Canvas), documented it, and added it to the HTML5 specification in the spirit of documenting what already works. So now we have widely supported Drag and Drop. (You can see Hickson's account of the process here:

<http://ln.hixie.ch/?start=%201115899732&count=1>.)

Why not just use JavaScript for drag and drop? For on-page elements you certainly could, but the thing about the (somewhat brain-damaged)

HTML5-via-1999 DnD API is it lets you drag and drop all kinds of content to and from *other applications*. There's a basic tutorial here: <http://html5doctor.com/native-drag-and-drop/>.

## Wrapping Up

So there you have it—a quick run through HTML5's web app oriented features. As you can see, HTML5 is about introducing a native web apps platform. Which isn't surprising considering the spec started as Web Applications 1.0 and Web Forms 2.0.

HTML5 is also just a stepping stone in the evolution of web applications; albeit a very important one, and one that has been a long time coming. Post-HTML5 developments (especially from Mozilla in mobile, as we discussed earlier) are continuing at a break-neck pace. They'll continue to be called HTML5, because as (I believe) Mark Pilgrim put it:

*HTML5 will always be popular, because anything that's popular will be called HTML5.*

It's an exciting time for the web. Hold on tight; it's going to be one hell of a ride.

# THE TRUTH ABOUT THE FUTURE OF WEB DESIGN: PERFORMANCE BASED DESIGN

I want to wrap up this book with a look into the future—not of HTML5, CSS3 or any particular technology, but the future of our profession. (And I want to pitch you on my upcoming web app at <http://itsninja.com>.)

Let's face it: we're nerds. We love technology. It's exciting, it's changing rapidly, and it's just plain fun to be on the forefront of one of the biggest technological and social phenomena in generations—the web.

But cool technology is a means to an end. It irritates me no end when web designers and developers breathlessly proclaim a new web site is “Designed with HTML5!” as though it means something. Technology *enables* design, but more/newer technology doesn't mean *better* design—sometimes it's quite the opposite. At the end of the day, web design is pretty simple. Users click or they leave. They engage or they bounce. They buy or they abandon. And it's the actual page design (and copy) that determines how often that happens.

So far, so obvious. But here's the kicker: *we can measure what users do with our design*. We can measure whether they click or buy, or bounce or exit. It's probably the most profound difference that separates our practice of designing on the web from just about any other. We can measure *performance* in a way no other discipline has been able to in the history of design.

As the saying goes, what gets measured gets improved. And how:

- 37signals A/B tested the design of their Highrise marketing site and improved conversions 102.5%: <http://37signals.com/svn/posts/2991-behind-the-scenes-ab-testing-part-3-final>
- Conversion Rate Experts applied their methodology (which they explain in detail) to a SEOMoz landing page and made them a cool additional \$1 million/year: <http://www.conversion-rate-experts.com/seomoz-case-study/>.
- Digital Telepathy (one design agency that really seems to get it) redesigned the CrazyEgg marketing site and improved their conversion rate 21%: <http://www.dtelepathy.com/case-studies/crazyegg>.

(That's just a taste, there are plenty more examples here: <http://abtests.com/> and here: <http://visualwebsiteoptimizer.com/case-studies.php>.)

## Operating In The Dark

It's just as well we can measure what we do, because right now we're surgeons operating in the dark, as we usually don't have a clue whether we help, hurt, or do nothing for our patients. We're operating on sites oblivious to design performance—whether people are reading more, or buying more, or bouncing less. We're not only operating in the dark, we're doing so while experimenting with a bunch of crazy techniques we've only just dreamt up and we're all very excited by.

That's a scary thought.

Doctors in the 18th century once thought grubby hands was a sign of professionalism, not a gross lack of sanitation. They thought they were doing the right thing. (The guy who tried to tell them otherwise, Ignaz Semmelweis, went mad—literally.) But when they started observing and measuring what was happening to patients, they figured out it wasn't such a good thing. Who knows what weird “best practices” we have in our profession that may turn out to be harmful?

## Performance Vs Production

It's not all doom and gloom, though. The beauty of measuring what we do is that we can objectively find the best version of *any* given design. Don't you hate coming up with a bunch of cool designs, only to have the client choose the one you like the least? (Or worse, something you're fairly sure will harm their business.) Wouldn't it be better if we could stop them from pulling the trigger when they point the gun at their foot?

We can do it now, and we can do it objectively, by changing the way we look at web design. I call it "Performance Based Design", and I think it's the next chapter of web design after "Standards Based Design" (long may it continue!).

Standards Based Design is about *how* we can achieve certain designs. It's production, and it's important. Much of this book has looked at how we can improve what we do when building sites (for example: use ARIA landmarks, use the new audio/video elements, use some new form features, experiment with Canvas and SVG, implement the History API, etc.). These are important developments in the *production* side of what we do.

Now, however, it's time to *also* start thinking about *what performs best* for our users. That is, what makes a real, measurable difference to how users interact with our site.

## Measure When You Redesign

I imagine you've read this book because you're going to be rolling out HTML5 site features, or whole site redesigns, as part of your day job. Maybe you're going to use more CSS3 too. And maybe you've been following the Responsive Web Design tsunami and are about to put out a hot new responsive version of your site.

If that's you, please measure what happens, and share the results!

Let's say you launch a "HTML5" site. It's fast thanks to the History API. There's some clever animation with SVG (or Canvas, or jQuery, or CSS3, or whatever). The video on the home page now uses a HTML5 media player.

What do you guess will happen?

Will bounce rates decrease? Does time on site improve? Do more people convert, or buy? We can measure all these things. If they do improve, great! We, as a community, need to know. We need data on what actually makes a difference for users—and what doesn't—so we can all learn from the evidence, not ideas or guesses or hopes or assumptions or “best practices”.

We have the data, we just need to start sharing it.

The same goes for responsive sites too. If you roll out a responsive mobile version of your site, what do your mobile users do? Do they bounce less, stay longer, and read more? Or does the opposite happen? Or nothing at all? Do you know how to find out?

Or maybe it's a tablet site. Does a responsive tablet design make *any measurable difference* to user performance? And if so, which designs work best? Simpler or more complex? Desktop-like or mobile-like?

There are so many questions, and guess what? *We already have the answers.* They are sitting in your Google Analytics account. We just have to dig them up, and share our data so we can learn what makes a real difference for users and what doesn't.

I've actually written a couple of books about this very topic, and how to integrate these concepts into your workflow. They're sitting unpublished on my hard drive at the moment, and I'm keen to know if you'd like me to put them out, so please let me know ([luke@itsninja.com](mailto:luke@itsninja.com) or tweet me [@lukestevens](https://twitter.com/lukestevens)).

## Let's Get Objective

This problem of designers (including myself) not having the data front and center when we redesign has been bothering me so much I've actually developed a web app that goes some way to solving this problem. Ninja for Google Analytics dives into your Google Analytics data and bubbles up the most relevant performance stats for you in a simple, elegant interface:

<http://itsninja.com>. Check it out, I think you (and your clients) will like it.

Objectively measuring design performance needs to become the *number one priority* for every one of our projects. It's bigger than HTML5 or any other technology floating around (as interesting as they are). When you start thinking in terms of measurable performance (conversion rates and engagement rates especially), you'll see web design and development in a whole new way.

Until then, go nuts with the new stuff in HTML5, measure what happens, and publish the results!

Thanks for reading.

Luke Stevens

<http://itsninja.com>

[luke@itsninja.com](mailto:luke@itsninja.com)

<http://twitter.com/lukestevens>

