

EXHIBIT 5



Smaller is Faster (and Safer Too)

Wednesday, July 15, 2009

We have just started using a new compression algorithm called Courgette to make Google Chrome updates small.

We have built Google Chrome to address multiple factors that affect browser security. One of the pillars of our approach is to keep the software up to date, so we push out updates to Google Chrome fairly regularly. On the stable channel these are mainly security bug fixes, but the updates are more adventurous and numerous on developer channel.

It is an anathema to us to push out a whole new 10MB update to give you a ten line security fix. We want smaller updates because it narrows the window of vulnerability. If the update is a tenth of the size, we can push ten times as many per unit of bandwidth. We have enough users that this means more users will be protected earlier. A secondary benefit is that a smaller update will work better for users who don't have great connectivity.

Rather than push put a whole new 10MB update, we send out a diff that takes the previous version of Google Chrome and generates the new version. We tried several binary diff algorithms and have been using bsdiff up until now. We are big fans of bsdiff - it is small and worked better than anything else we tried.

But bsdiff was still producing diffs that were bigger than we felt were necessary. So we wrote a new diff algorithm that knows more about the kind of data we are pushing - large files containing compiled executables. Here are the sizes for the recent 190.1->190.4 update on the developer channel:

- **Full update:** 10,385,920 bytes
- **bsdiff update:** 704,512 bytes
- **Courgette update:** 78,848 bytes

The small size in combination with Google Chrome's silent update means we can update as often as necessary to keep users safe.

More information on how Courgette works can be found [here](#).

Posted by Stephen Adams, Software Engineer

50 comments:

[Joel Martinez](#) said...

Impressive ... thanks for talking about this. Hopefully others will take note and implement similar mechanisms. This will only make the entire ecosystem better :-)

July 15, 2009 12:46 PM

[Drew](#) said...

I love you Google.


...chromium.org/.../smaller-is-faster-an...

Search our Blog

Archive

November (1)

Subscribe

 [RSS Feed](#)

More Blogs from Google

Visit our [directory](#) for more information about Google blogs.

Useful links

[Chromium Homepage](#)

[Google Chrome](#)

[Google Chrome Release Notes](#)

[Google Open Source Blog](#)

[Google Code Blog](#)

[WebKit Blog](#)



July 15, 2009 1:13 PM

Don-J-Rude said...

Great idea. Move all the (simple) pointer recalculations onto the client to perform after the diff of the "body" is applied by using a disassembler! The hinting also makes total sense, use some standard magic bytes to identify binary types and work in the decompressed format.

Memory usage analysis? CPU usage analysis? Perhaps the future paper will explain but it seems very CPU intensive (common trade-off).

July 15, 2009 1:13 PM

Jan de Mooij said...

Very nice, this inventiveness is what i love about Google.

Could be very useful for other vendors (Mozilla etc) as well...

July 15, 2009 1:51 PM

Philip Arnason said...

I'm continually impressed with google's strive for speed. I've recently started using chrome almost exclusively. It's unreal how long it takes for a new tab to open up in MSIE.

July 15, 2009 2:09 PM

John said...

Have you looked into my binary delta algorithm, BDelta? The GPL'd source code is available at deltup.org, and my tests show that it is much faster, much less memory-hungry, and produces better deltas than bsdiff. The released code is old, but still very usable, and I have some updates in the pipeline. Just note that the output of the program is not compressed, so you'll want to pass it to a good compression tool like bzip2.

If you're interested, I'd be glad to help you with the implementation. If not, why not? ;-)

July 15, 2009 2:16 PM

greenPhil said...

1. Versioning is very important in this case.
2. One diff can go to all the clients with same version.
3. You don't need to calculate diff for each installed chrome browser.
4. Just calculate the diff(say version x.y x.y+1) in the upgrade server and push it to all browsers that have x.y version.

This can work very well if there is no change in format files that can potentially be touched by the browser user.

July 15, 2009 2:19 PM

Kent Fenwick said...

That is so cool, you read about that stuff in Algorithms but to see it in Practice is really cool.

Its the little things that make you so impressive.

Kent

July 15, 2009 3:01 PM

Harold Fowler said...

11/19/2009

Chromium Blog: Smaller is Faster (and ...

Hey, if its safer I am all for it dude!

RT

<http://www.anonymize.tk>

July 15, 2009 3:38 PM

gmx said...

Sweet

The bad part is the "silent" update.

July 15, 2009 3:56 PM

jack.feed said...

Firefox needs to make their updates smaller like Chromium. Bravo!

July 15, 2009 6:25 PM

rdza said...

This post makes "ChromeOS smartphone" ring in my head

July 15, 2009 9:03 PM

B.D. said...

This post has been removed by the author.

July 15, 2009 9:53 PM

BrainAds said...

still can't believe that chrome is only for windows users, how about fixing that in an update?

July 15, 2009 10:12 PM

muharem said...

Please make the source code available under GPL (or similar) so other projects/people can use it.

July 15, 2009 10:53 PM

Ralph said...

The Courgette algorithm sounds tricky to test and debug. Was it a long process to get it working properly?

July 15, 2009 11:00 PM

Sankar said...

Isn't the same technique already available as delta-rpms ?

July 15, 2009 11:41 PM

Jaco Pretorius said...

Sounds like a real pain to develop and debug - well done

July 16, 2009 12:05 AM

Papa Marteau said...

Wondering why any vendor think about it before. Probably the diff between versions is difficult to maintain. Especially if your soft is not a web browser.

July 16, 2009 12:44 AM

Guangrong86 said...

...chromium.org/.../smaller-is-faster-an...

[url=http://www.brighttrans.com/en/index.php]Translation company[/url]
thanks for talking about this, hopefully others will take note and implement similar mechanisms. this will only make the entire ecosystem better.

July 16, 2009 12:59 AM

ahsankhan said...

If you can sell a new car for 5000 then your dealer network will grow faster than you would believe. motor trade recruitment

July 16, 2009 1:47 AM

vincent b. said...

You seriously called it Courgette?
:-D

July 16, 2009 1:55 AM

Fabio Turati said...

Is this the reason that last week there were many more updates (3) than usual? Does it mean we can expect more frequent updates from now on?

Anyway, great job guys, the compression you achieved is impressive.

July 16, 2009 4:13 AM

Michel En La Red said...

Realmente Chrome es una super maquina, felicitaciones a Google, y a todos los colaboradores de este excelente proyecto.....

July 16, 2009 4:21 AM

Art Vandalay said...

Impressive, although calling it a "compression algorithm" could be misleading for people who don't bother to read and understand the details. Such people might assume "oh, look, a stupendously better general algorithm for shrinking binary data", which is not at all the thrust of this clever idea for reducing the necessary size of *binary diffs of object code*.

July 16, 2009 6:27 AM

famille reuel said...

Impressive.

Is it just that everyone else is lazy or stupid, or that Google simply refuses to not question to status quo?

Once again, simply well done and impressive.

July 16, 2009 7:28 AM

Óscar said...

Will it be released as free software?

July 16, 2009 9:44 AM

PWRAdmin said...

Will you guys be open sourcing this amazing little tool!

July 16, 2009 9:52 AM

Balaji said...

sounds cool . is it available for 3rd parties to use it???

July 16, 2009 10:20 AM

Michel S. said...

@Sankar

Probably not. The deltarpn package that Fedora ships, for instance, does not even do binary diffing -- it relies on bzip2. With delta RPMs, the idea is that most files remain unchanged between two different versions, and you can just transfer the changes. Within each file, it's just normal compression.

July 16, 2009 10:47 AM

Greg said...

go google!

July 16, 2009 12:30 PM

Stephen said...

This post has been removed by the author.

July 16, 2009 5:31 PM

Stephen said...

@John-of-BDelta

Email me at sra@chromium.org so we can figure out how to get you some sample input files for an apples-to-apples comparison.

@muharem, @Oscar, @PWRadmin, @Balaji
Courgette is open source as part of the Chromium project. See <http://code.google.com/chromium/terms.html>

July 16, 2009 5:40 PM

Nicolás said...

Excellent stuff. If only more companies followed your philosophy.

July 16, 2009 7:15 PM

ssj4Gogeta said...

Another example which shows that Google is made of "geeks" and not "suits".

July 17, 2009 1:05 AM

Ralph said...

You may like to know that Colin Percival, bsdiffs author, recently made a request that if BigCo uses some FLOSS they offer one of the developers some schwag, e.g. a T-shirt.

<http://www.daemonology.net/blog/2009-07-14-a-call-for-schwag.html>

His tarsnap "backup to the cloud" system is shaping up nicely by the way.
<http://www.tarsnap.com/>

July 17, 2009 4:27 AM

Jonathan said...

re: Michel S.
deltarpn does do a binary diff (using a modified bsdiff algorithm, with the diff stored using bzip2), but it also uses an algorithm described in <http://www.daemonology.net/papers/thesis.pdf> to deal with the same problem
courgette is trying to address (though in a far more generic way). I would be very interested in generating a diff using the deltarpn algorithm and the courgette algorithm and see what their respective savings are.

11/19/2009

Chromium Blog: Smaller is Faster (and ...

July 18, 2009 5:22 AM

Jan Kotek said...

You did not make new diff algorithm. It is ugly hack to the bsdiff for one very specific case.

July 18, 2009 8:11 AM

Jeet said...

Google are Great Forever.... We Love Google....

July 20, 2009 12:20 AM

conrad said...

I don't quite see how this is a new "compression" algo. You're making the diff's smarter and that is what is making the patches so much smaller.

July 20, 2009 11:22 AM

user said...

gucci mens shoes site: <http://www.guccimensshoes.net>

Replica Louboutin site: <http://www.repicalouboutin.com>

hugo boss inmotion site: <http://www.hugobossinmotion.com>

replica louis site: <http://www.repicalouis.com>

nokian99 site: <http://www.nokian99.net>

edhardy t shirts site: <http://www.edhardytshirts.net>

SEO site: <http://www.seo9seo.com>

Wholesale site: <http://www.ourwholesale.net>

B2B site: <http://www.b2btradeplatform.com>

July 21, 2009 12:45 AM

Sameer said...

Makes total sense given the more frequent nature of updates on the dev channel.

July 21, 2009 8:15 AM

عارف بن ياسوكي said...

Cool!! can't wait the opportunity to use this in my project!

Afriza N. Arief

August 4, 2009 1:14 AM

cheguaka said...

Fedora has chosen XZ instead of bzip2 for rpm compression. Worth to check?

<http://tukaani.org/xz/>

August 6, 2009 3:19 AM

Kevin said...

This really is impressive - its innovation like this that makes chrome feel more cutting-edge than other browsers.

August 18, 2009 2:46 AM

Eldad said...

Hello,

Chrome is GREAT, thanks very much, you are doing a great job guys (and gals)!

It is now my default browser, taking over from the excellent but slower Avant Browser.

A couple suggestions: Please add (perhaps to "Control current page"): "Email page",

...chromium.org/.../smaller-is-faster-an...

6/7

GOOG-00026264

11/19/2009

Chromium Blog: Smaller is Faster (and ...

"Email URL", "Email selected", and if possible "Print selected". (selected text and/or graphics).

Avant Browser has a great feature - Forward and Back pages by: Right click then Left mouse click to forward, Left click then Right click to back a page. It is a great feature that really speeds up surfing. I hope you can add that to Chrome.

Thanks again,
Eldad (Saugerties NY)

August 22, 2009 12:34 PM

raju said...

this is a good hacl dr. adams.
please try to embed a scheme interpreter into the next chrome update.

August 28, 2009 9:11 PM

ikpi! said...

Good!! So, How to download the source code?

September 15, 2009 2:15 AM

Christophe said...

Stephen,

I would love to hear your word on my proposal - please see here:

http://groups.google.com/group/chromium-discuss/browse_thread/thread/eb9ee72b32cadaef

I created a PDF explaining the process which can be downloaded here:

<http://www.multimedial.de/share/BitPatternLibOverview.pdf>

I am looking into doing a reference implementation in python of this.

I would very much appreciate any kind of input you got to this. So far, I came to the conclusion that this is a modified Huffman encoding scheme with a pre-build and non-optimized dictionary.

Any input very much welcome,
Christophe Leske

September 30, 2009 2:08 AM

vtorlazo said...

is the source code available somewhere? ca we see the implementation?

November 14, 2009 10:35 AM

[Post a Comment](#)

[Newer Post](#) [Oider Post](#) [Home](#)

Subscribe to: [Post Comments \(Atom\)](#)

©2008 Google - [Privacy Policy](#) - [Terms of Service](#)