

<http://j.mp/ccog2015>

# Data Pipes & Relevance

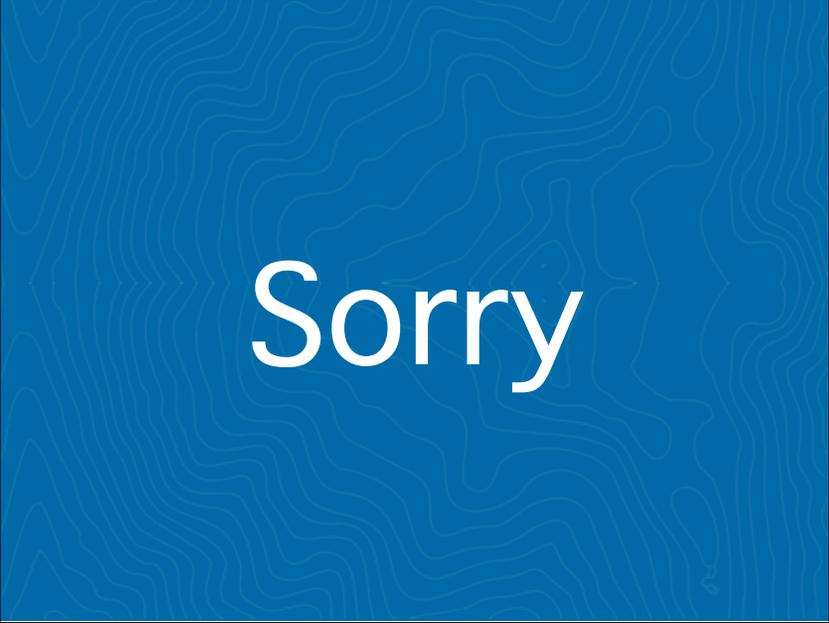
What can be done for government mapping



<http://j.mp/ccog2015>

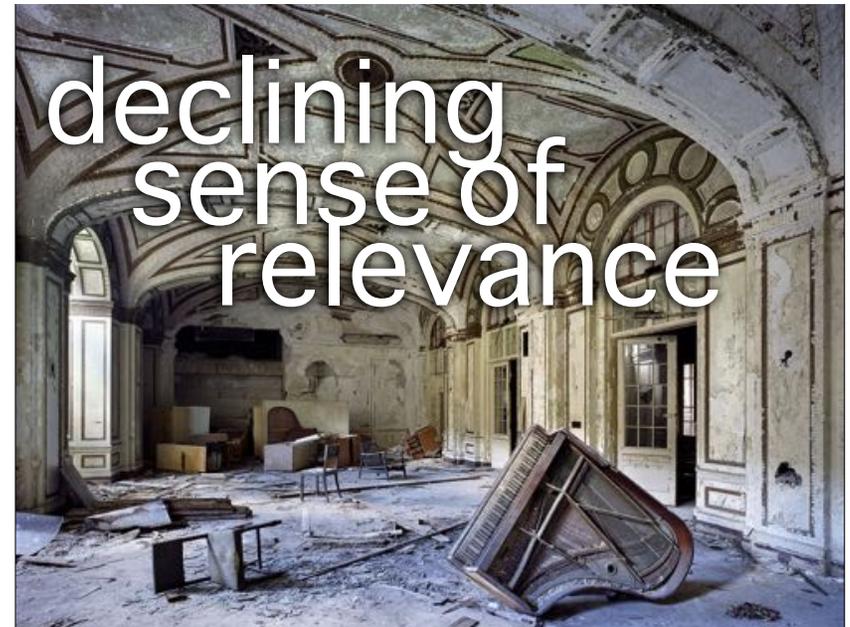
Greetings from start-up land.  
I work for a spatial visualization and analytics start-up, CartoDB, now, but I'm actually a refugee from the world of government mapping. I spent 10 years directly consulting to BC and Canadian mapping agencies here in Victoria, and the past six years working indirectly for US defense intelligence and civilian government mapping agencies as part of an open source support company. So, the last twelve months in start-up world have been very A-typical of my career, in that I haven't touched a government use case once.

I'm going to be making reference to a lot of different sites, and articles and talks, and rather than go into them in detail, or having you transcribe references one at a time, I've put all the links into a reading list online, so if you're interested you can peruse them later at your leisure,



# Sorry

I didn't really know how approach this talk,  
and given what I came up with,  
I kind of feel like I should start with an apology,  
So, I'm sorry,  
<X> if I say anything hurtful or overly critical,  
because any criticisms I make of government geomatics  
come out of love, not anything else,  
because I have worked with government mapping for  
a very long time, first with what was  
once known as geographic data bc here, and later with  
natural resources canada and more recently with the  
national geospatial intelligence agency, and what these  
one proud agencies have in common

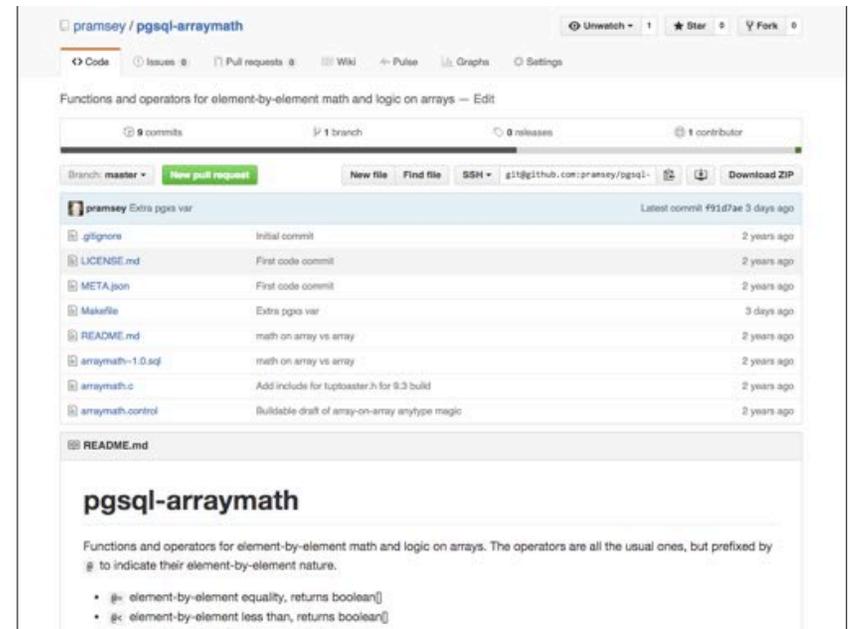


is a declining  
sense of relevance in the current era,  
a decline that has been brought on by huge  
external  
shifts in the industry, and in the technology of  
location and place, and in the management and  
processing  
of data, exacerbated by institutional inertia and a  
very human tendency to place a premium on How  
Things Have Always Been Done, and a tendency to  
extrapolate What Is To Come from What has  
Already Happened

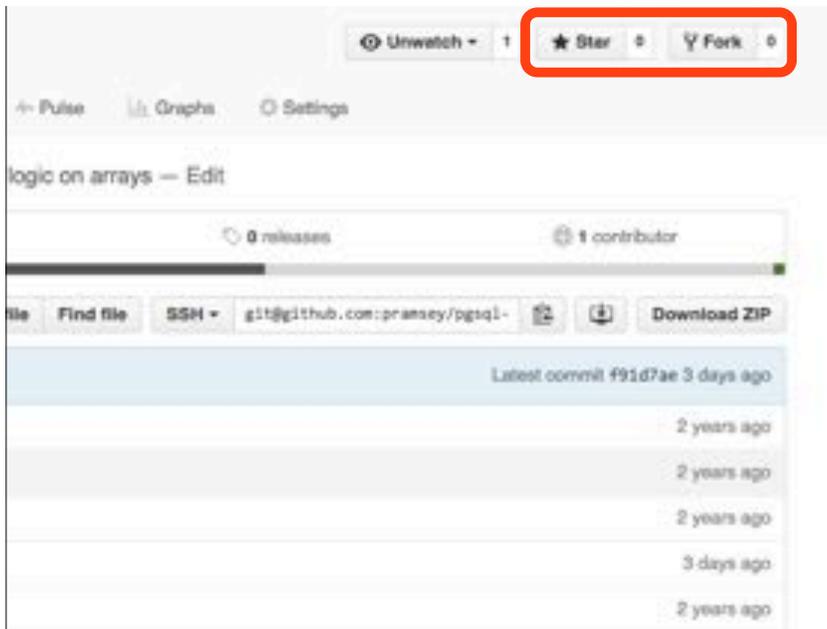
# Relevance

When I say that the relevance of government mapping has declined, I mean relevance in a very broad, yet important sense.

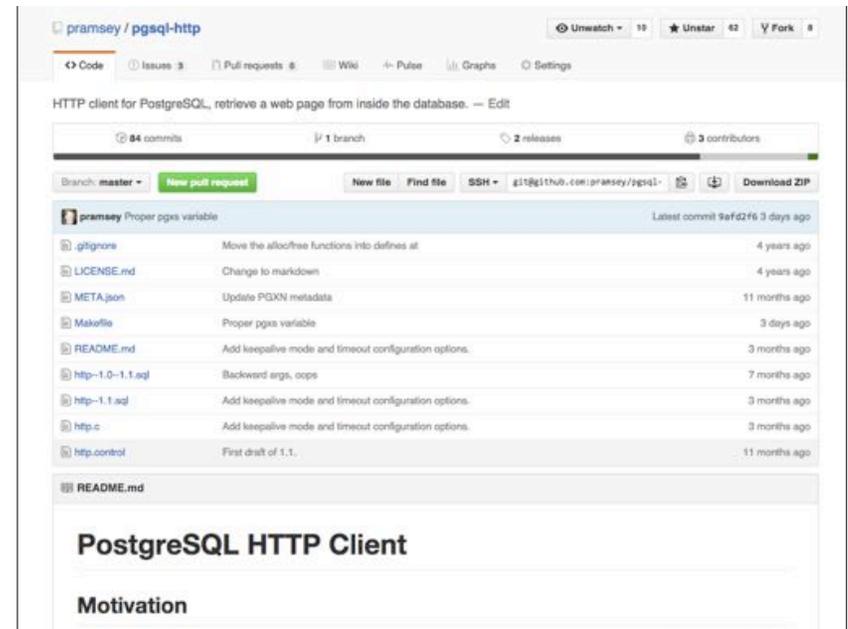
I have made my world for the last 10 years as an open source software developer, so I have a pretty keen sense of what it means to be relevant, and how relevance as measured by the cruel world can differ from relevance as measured by my internal sense of good and bad.



I have lately taken to publishing small extensions to the PostgreSQL database, and one of my latest is an extension that does element by element math transforms on arrays. It's really really cool, it handles any operation that makes sense for the array element types, but without my having had to code any of that, it just binds into the internal system catalogs of the database, so it's simultaneously a very small piece of code that provides a very wide swath of functionality.



Unfortunately, nobody really cares. It lacks relevance.



On the other hand, before that I did an extension that allowed any web address to be called from a function and returned the content from the address. It's got a certain brutal utility, but it's just one huge function running HTTP calls, and it has some critical shortcomings in terms of blocking the query thread until the HTTP request returns.

## PostgreSQL HTTP Client

### Motivation

Wouldn't it be nice to be able to write a trigger that called a web service? Either to get back a result, or to poke that service into refreshing itself against the new state of the database?

This extension is for that.

### Why This is a Bad Idea

- "What happens if the web page takes a long time to return?" Your SQL call will just wait there until it does. Make sure your web service fails fast.
- "What if the web page returns junk?" Your SQL call will have to test for junk before doing anything with the payload.
- "What if the web page never returns?" I've found this code can really hang a back-end hard. The curl timeout settings need more testing and tweaking for faster failure and timeout.

### To Do

- The new <http://www.postgresql.org/docs/9.3/static/bgworker.html> background worker support could be used to set up an HTTP request queue, so that pgsqj-http can register a request and callback and then return immediately.
- Inevitably some web server will return gzip content (Content-Encoding) without being asked for it. Handling that gracefully would be good.

I even included a section in the documentation called "why this is a bad idea".

The screenshot shows the GitHub repository page for 'prasey/pgsql-http'. At the top right, there are three buttons: 'Unwatch' (10), 'Unstar' (62), and 'Fork' (8). The 'Unstar' and 'Fork' buttons are highlighted with a red box. Below these buttons, there are links for 'Pulse', 'Graphs', and 'Settings'. The main content area shows the repository name 'prasey/pgsql-http' and the text 'the database. — Edit'. There are also links for '2 releases' and '3 contributors'. At the bottom, there are buttons for 'Find file', 'SSH', and 'Download ZIP'. The latest commit is '9efd2f6 3 days ago'.

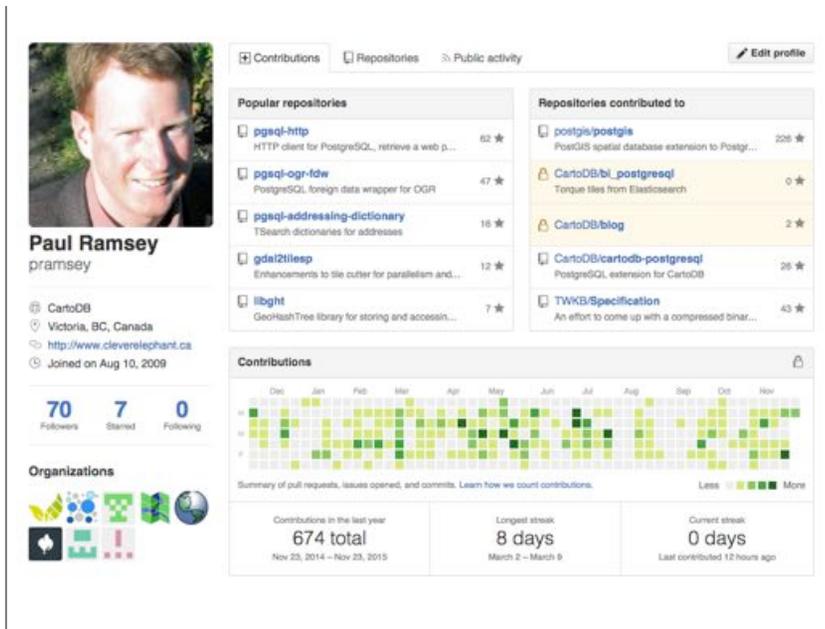
Nonetheless, it's really really popular. It has relevance.



It probably sounds very millennial of me to equate popularity with relevance, but in my defense, I'm actually a Gen X'er, and this is very much how the world works now. We recently had an election, or shall we say,



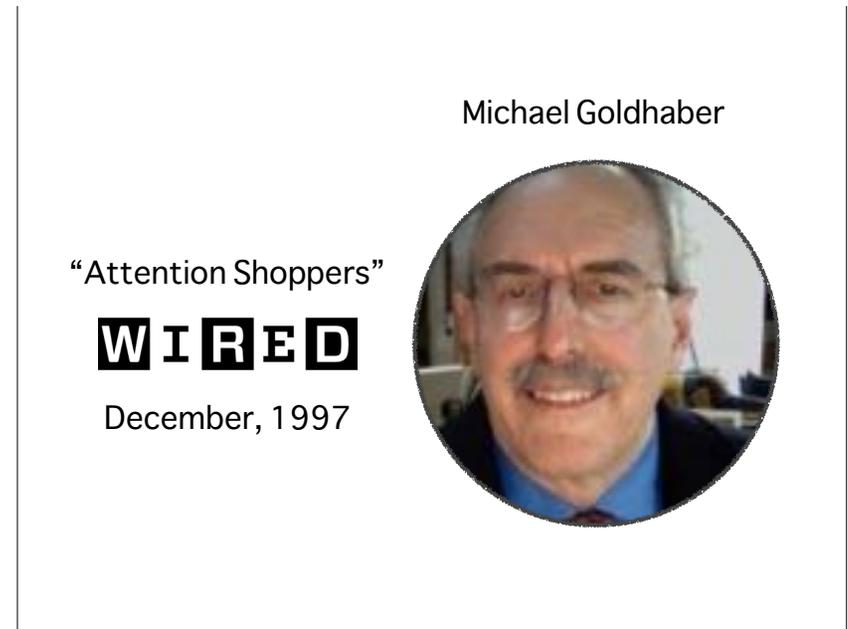
a relevance contest, and the most relevant party now holds the reins of power. I know as an open source developer, and speaker, my future earnings power depends very much on my ability to



create and maintain projects that are relevant to enough people that I can earn a living from the ones who need developer services.

Working on what I feel like, without consulting their overall relevance, would be a recipe for eventual penury.

Michael Goldhaber wrote an article about this 20 years ago, at the dawn of the internet,



and I remember reading it in Wired magazine at the time, and thinking "this is the stupidest, most improbable thing I've ever read", and yet now here we are 20 years later and it turns out he predicted the dynamic of content and intellectual capital on the internet extremely well.

What he described was an "attention economy", where the most important value a person or organization can accrue is attention. When people grant you their time, they increase your relevance. The Bank of Canada can print more dollars, but they cannot print you more time, you have a fixed amount, and when you provide it to someone, you're gifting them something of very high value.



So the very hard math, in the internet era, is to the extent that you're garnering attention, you have relevance.

For content, that means that the new optimal price point is free, to gather up the most attention available, with the least friction possible.



I mentioned earlier the long decline of government mapping, which can be described in terms of a loss in in the share of the publics attention over time.

Things were better in the period of paper mapping,



when the market for map information was much smaller, but in which the products of government mapping agencies formed the basis for most mapping products, in one way or another, from maritime charts to topographic maps, to the basic backing data of street and road atlases.



When I got involved in mapping in BC, in the early 1990s, Geographic Data BC still had a staff of over 100, though the glory days of the organization were already a decade in the past. At that time, most maps the public would see were still based on government data.



Similarly, in the glory days of the national geospatial intelligence agency, (when it was still the NIMA) the whole world was mapped to NIMA standards, and the source of the finest, most complete maps of the entire world was NIMA. The admirals and generals didn't have anything else to compare it to, and nothing else could compare.

$$y = x \cdot \text{change}$$

Things are pretty different now.

(nerd joke... get it? the only constant is change)



The largest national mapping agency in the world is headquartered in Mountain View, California, in a building decorated like a hypertrophic kindergarten.



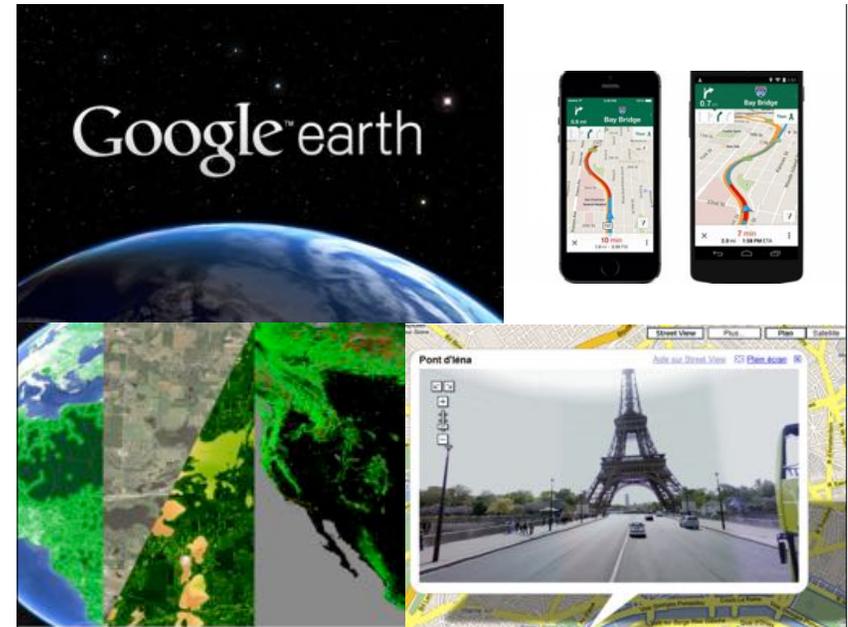
"Google has 7,100 people working on Google Maps. That's 1,100 full-time employees and 6,000 contractors who drive street view cars, fly planes, and correct map details."  
--wired <http://www.wired.co.uk/news/archive/2014-12/09/google-maps>

And those 7100 people aren't working in the old, manual processes of the last generation, they are taking advantage of the finest machine learning and computational efficiencies the world has to offer. And they aren't just collating other people's data anymore, as they were in the early years.



They have massive primary data capture going on: they commission their own imagery acquisition, they drive their famous fleet of street view cars on every street in the world it's legal, they even have their own orbital sensors now, since purchasing Skybox Imaging, with two in satellites orbit and six more to be launched next year.

<http://www.bloomberg.com/news/articles/2014-06-10/google-buying-satellite-company-skybox-imaging-for-500-million>  
<http://spacenews.com/vega-to-launch-skybox-satellites/>



Although the NGA still employs more people than Google, 14,500 employees and contractors, I have no doubt that Google outproduces them handily. Certainly the heirarchy of wonder in the Defence Department have been inverted. Instead of being awed by the information their reconnaissance arm brings to them, the Generals and Admirals bring the latest consumer wonders of Google and its ilk to the NGA, and say "why aren't you giving us something as good as this?"

relevance = wonder<sup>2</sup>

Which is perhaps a familiar story, to this crowd.

**Our New  
Master**

So, everyone is feeling the squeeze,  
because the geomatics world has a new master.  
Probably you all have one.



Do all of you have one of these?  
Regardless of how you enumerate your stakeholders,



there's no escaping the fact that **these things** are your  
real customers now.  
The taxpaying citizenry,  
the people government is supposed to serve,  
own these things at a staggering rate.



More importantly perhaps, so do the real bosses, our elected officials. And they are all consuming spatial information at an unprecedented rate, using these devices. Where a street map might have lived in a glove compartment and been consulted once a week,



digital maps are being consulted directly and indirectly on a minute-by-minute basis now.

That's the environment. We all know it. So if we believe that government mapping is supposed to be serving the public, the ur-question that should drive every strategic decision is

**“how do we get our data onto those devices, and into the pockets of our millions of citizens.”**

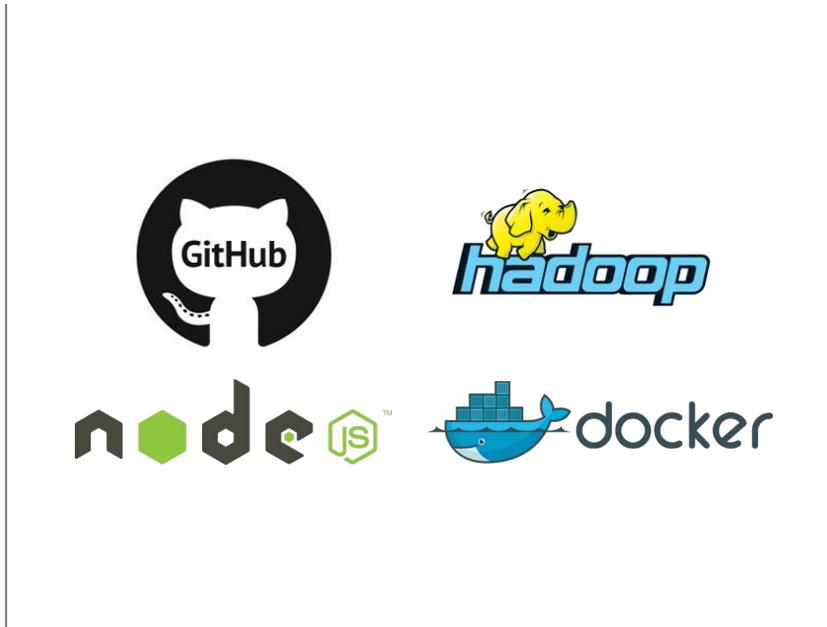
# Developers as Gatekeepers

The answer, to an extent, is via developers

In the software development world, marketers have long been aware that the gateway to category success is no longer via the CEO's office, but via the tech workers cubicle.



Basically, developers choose the tools they like first, and the enterprise eventually standardizes on them.



So, git and github are new standards,  
docker is increasingly becoming a containerization  
standard,  
node.js is the technology of choice for middlewares,  
and the surge of all these technologies has been driven  
bottom up.

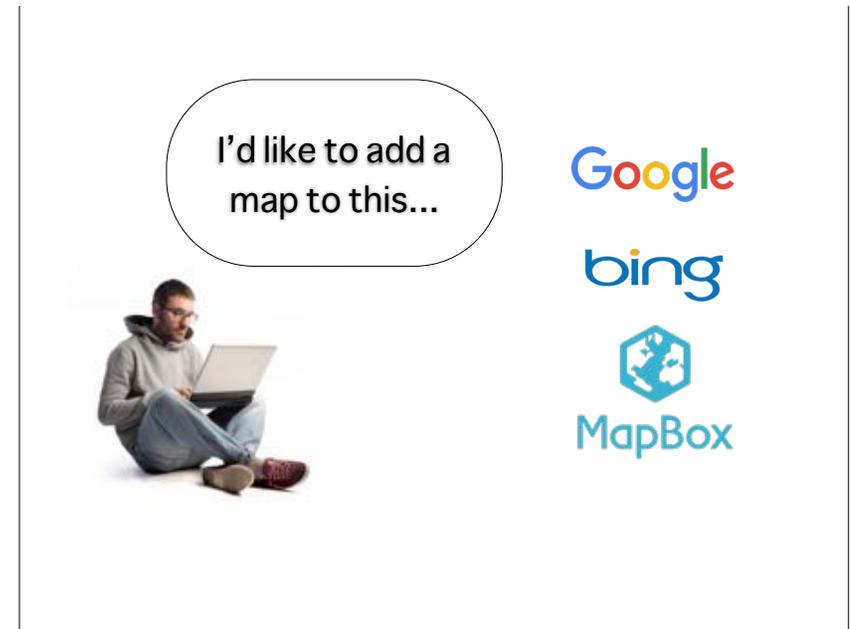


So, I have another story, which you might not be  
familiar with,  
but which is as worrying in its own way.  
From my perch, on the edge of start-up world,  
and the open source world,  
I can observe developers in their native habitat.

Mostly, start-up and open source developers  
know **bupkus** about mapping and spatial.



As recently as 5 years ago, there would be a 50% chance that a new developer trying to do something with spatial would gravitate to government sources of data. <X> Often US Census TIGER, as a large complete corpus, <X> and in Canada to GeoGratis and GeoBase. And for those people, my knowledge of other interesting sources of government data would always be charming, particularly in the days before open data portals, <X> when knowing where the US Department of Agriculture or Department of Transportation hid their data was a good party trick.



But not anymore. Developers starting fresh don't look at government data anymore. They get their dose of spatial in a couple ways:

- At the low end, if they are interested in consuming a little bit of location context and some services, they'll use an API provider like Google or Microsoft. If they want to avoid the personal data tracking inherent in the large services they'll use a smaller provider like MapBox.



- At the higher end, if they are doing something more bespoke and need to work from raw data, they'll get the overview data from Natural Earth, place data from GeoNames, and the detailed data from OpenStreetMap.

That's it. If you, as developer, can get **data of sufficient detail and scope to support your project from a single source**, you'll do so and then you'll **stop**.



Why look at Canadian government data, when (a) it only covers Canada (or worse, only one province in Canada) and (b) it's probably not as current as the other data sources.

It's hard to overstate ... ease of access and, ease of use in the decline of relevance of government mapping.

When people consumed mapping via **paper**, government data stood at the headwaters of the creation process for those products.

In the **digital age**, government's prime place in the data watershed has been usurped by community data sources and privately compiled data.

# Distribution

So why is that?

If you accept, to some extent, that relevance is conditioned on the number of people consuming your data, then the process of distribution, of getting into the pipeline for product creation, should loom large as a topic of interest.

There is a threshold moment, between the time when data is in **your custody** and control, and the time when **someone else** is using your data, and that moment is when distribution happens.

Here's a couple of examples of government taking a non-traditional approach to distribution,

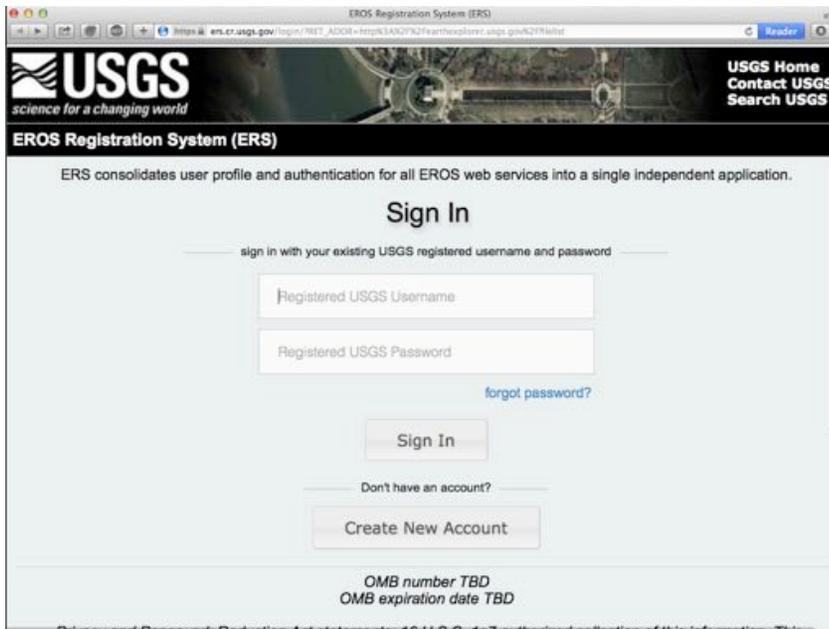


Last year a huge change happened in the distribution of LandSat data.

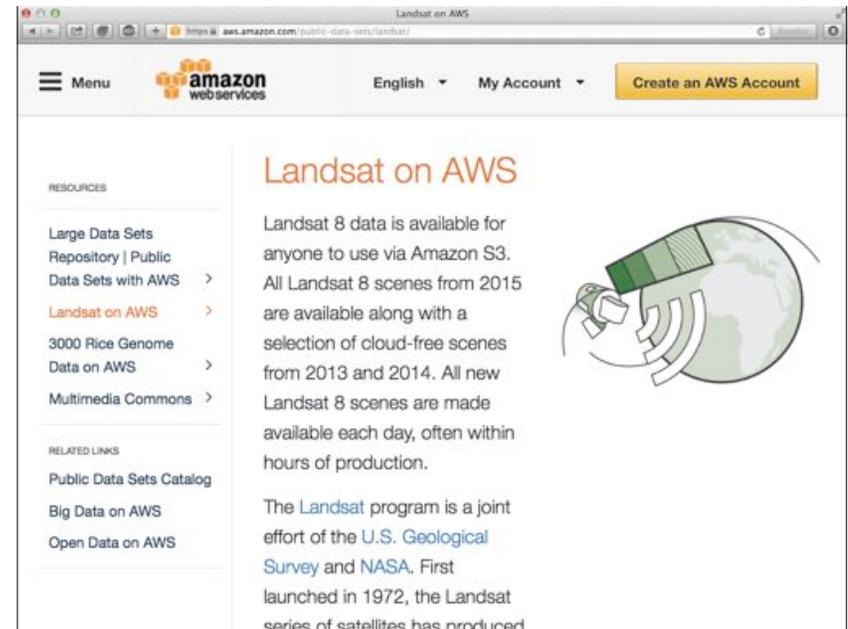
LandSat is one of NASA's longest running earth observation missions, covering multiple decades.

LandSat data has been publicly available for a long time.

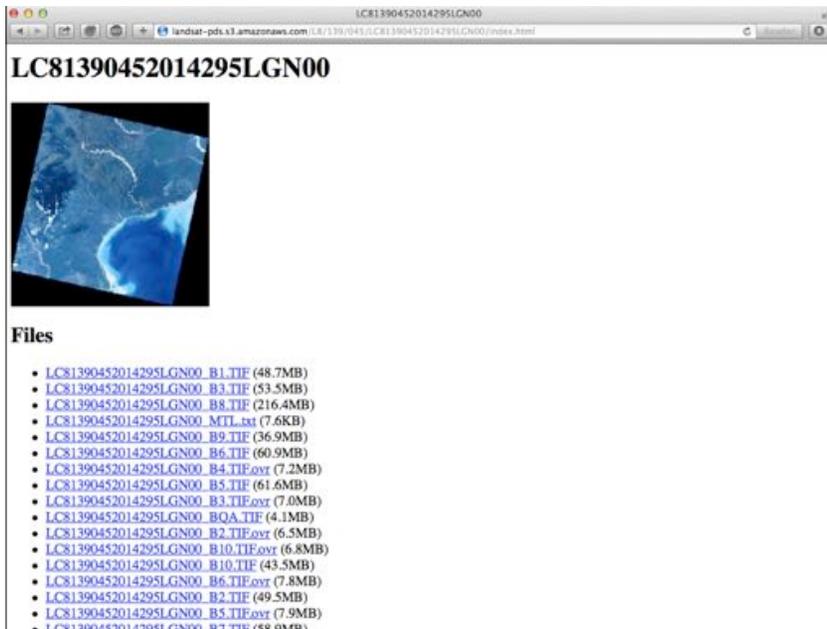
Because the sensor runs continuously, the data volume is very very large, which has made distribution hard.



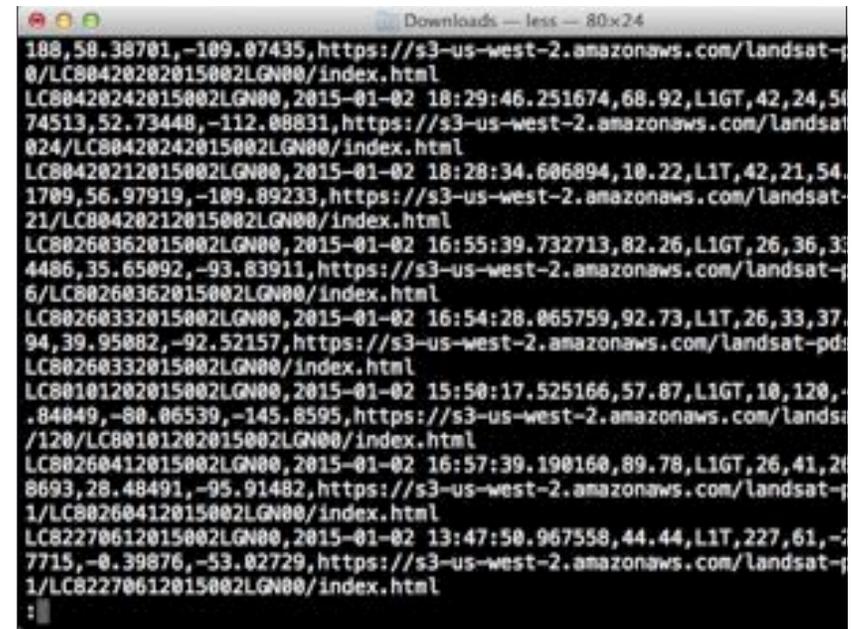
NASA also got stuck, in the internet age, into a very common pattern, of spending a lot of effort on trying to make distribution "easy", via portal and data selection software, without making **bulk access** fast and simple first.



The big change this year was the way LandSat 8 data is being distributed. It may still be available on an "easy" portal, but that's not what got the developer community excited. The big change that everyone got excited about was LandSat 8 data is now available in basically real-time on the AWS cloud.

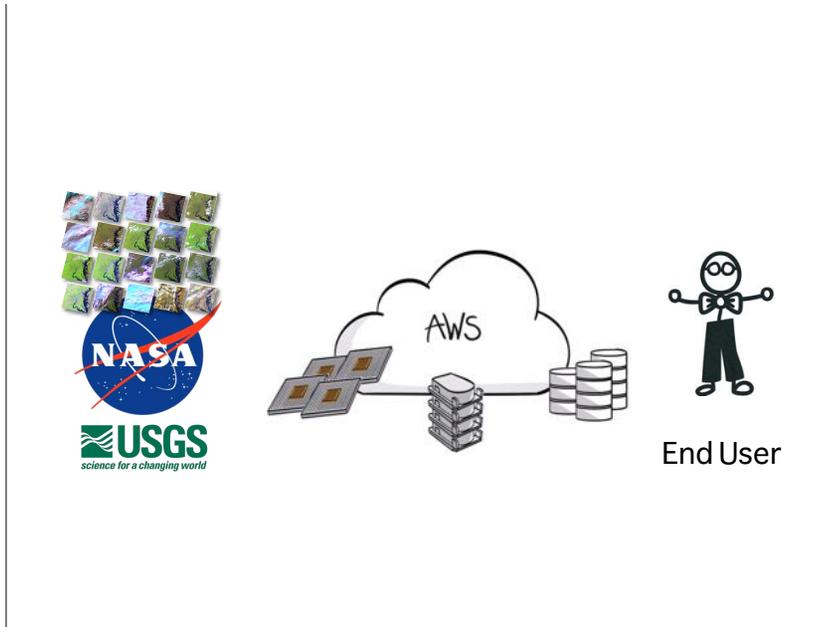


The distribution is not user friendly at all, but it's simple, and it's computer friendly. Each new raw scene gets stuffed into an S3 bucket in the AWS cloud,

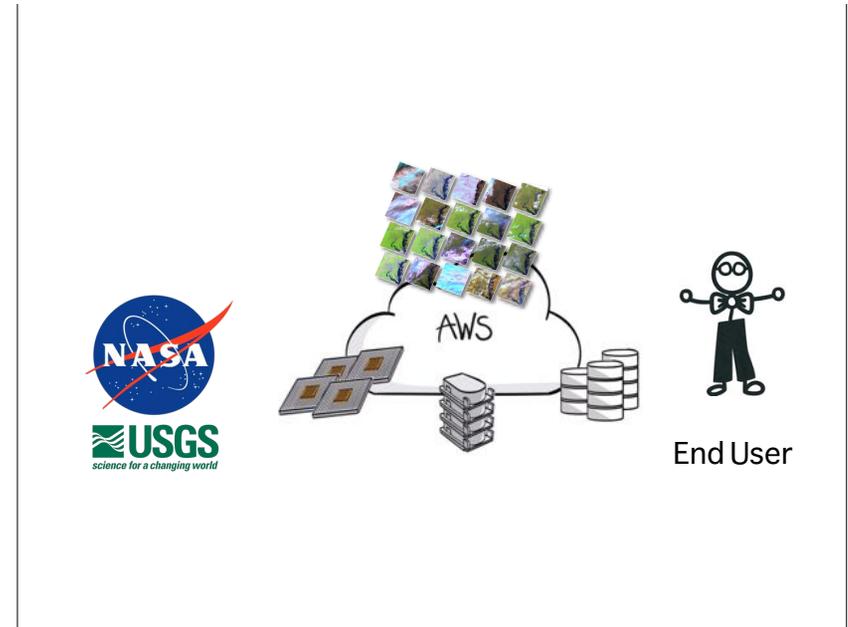


and the text metadata file is updated with information about the new scene. That's it. (And I got both the previous links to raw data in about 30 seconds off the AWS page.)

It's a big deal because it removes two problems with portal-based locally hosted distribution.



First,  
Putting the data into the cloud moves it very close to the point of consumption.  
And it's a fact of modern computing, that if you're doing high volume computation now, you're doing it in the cloud,

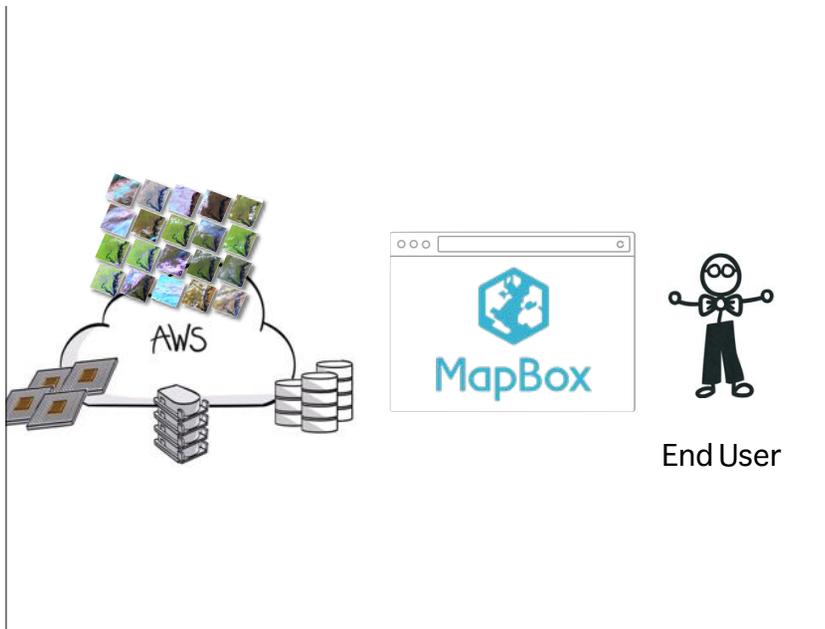


so having your data close to your processing speeds things up immensely.

I was reminded of this recently working on a big data problem at CartoDB, slinging multi-gigabyte files between servers.

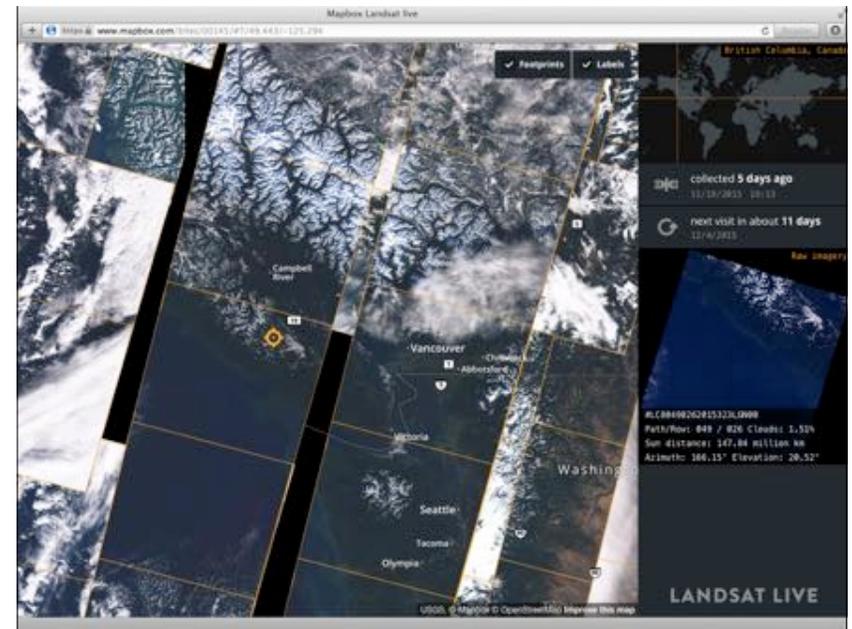
Ordinarily, hitting "copy" on a 4 gigabyte file is an opportunity to stand up and get a fresh cup of coffee, browse the sports section for a little while.

Moving data between servers in an Amazon data center barely affords time to touch your toes a couple times.



Second, Making the process simple, raw data in a bucket and metadata in a file, makes automating processes off the data stream exceedingly easy.

On the back of this real-time stream of imagery it was very easy for MapBox to build a value-added real-time stream of their own, taking the raw data, rectifying it, reprojecting it, and making it part of a real-time tile pyramid suitable for use in any internet mapping application.



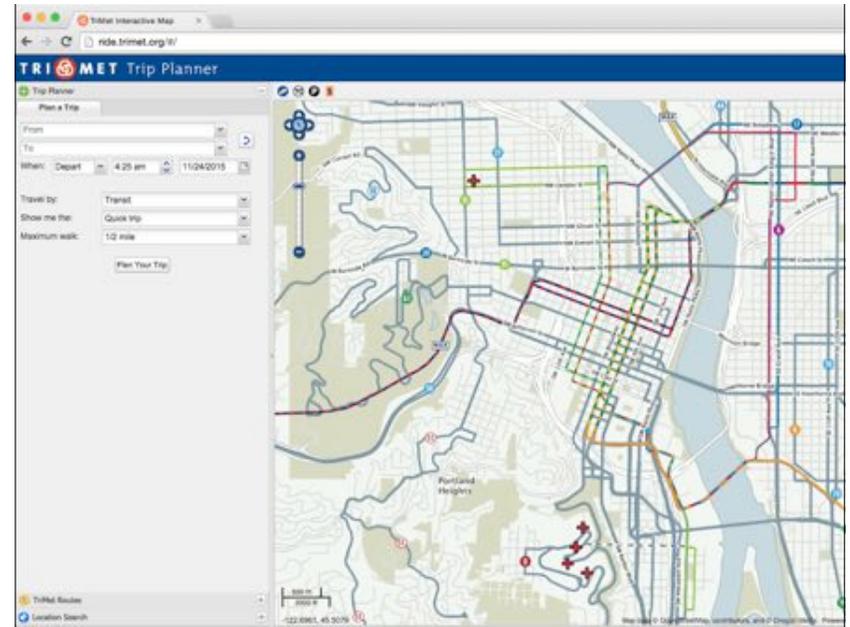
This kind of third party service would not be possible against a NASA-hosted portal. The transfer volume of multiple users pulling the data over the NASA pipes would add costs to NASA in proportion to the popularity of their data, and the complexity of portal access would both slow down external development.

The second non-traditional distribution example is exceptionally wierd,



Portland Trimet runs the public transit in the Portland Region.

To provide internal and external maps of their routes, and other spatial services, they need an up-to-date road network.



Functionally, that means they need to manage it themselves, so they can update it as reality changes in the Portland transportation system, roads are closed, rebuilt, and so on.

They had previously built an internal product by modifying a licensed copy of data from a road data vendor, but the annual data costs were high.

# Improve



They saw that an incremental investment in a public domain road data set could cut costs, but the question was: enhance TIGER and still have an internal data set or enhance OpenStreetMap. They chose to enhance open street map.



They hired two interns for a year to do all the changes to OSM they needed, at a cost less than their annual licensing fee for their old data, and when the major changes were done they moved the task to a part-time load on a regular staffer.



Effectively, there is a dataset, the "portland trimet roads data", which they maintain, but they now neither host it themselves, nor distribute it themselves  
<X> because their data is embedded in OpenStreetmap. And because it is embedded, their data is now used by all manner of third parties building maps in Portland.  
So, as the most active and professional map managers in the area, they've become the arbiters of truth for much of the road network people see every day in their apps and their maps.  
They've become **MORE** authoritative by abandoning the quest for control.

We *influence*  
~~control~~  
the data

It's a novel form of distribution and it turns the idea of data custodianship on its head.  
<X> It's similar in a lot of ways to open source.  
"My" project is the PostGIS spatial database, but it's only "mine" insofar as I'm a major contributor and thus guide the development and priorities of the project. I don't own the intellectual property, and yet, that somehow doesn't make the project any less mine, in the eyes of the world, or my own eyes.

# Open Street Map

I've mentioned OpenStreetMap a number of times, and I think it's worth elaborating a little bit on it, since I've found a certain level of hostility to OSM in data folks from time to time.

So let me quickly address some common concerns,

“It’s not  
authoritative”

"It's not authoritative"

On authoritative, I'll grant that, but will also note that,

Authoritative



Relevant



WIKIPEDIA  
The Free Encyclopedia

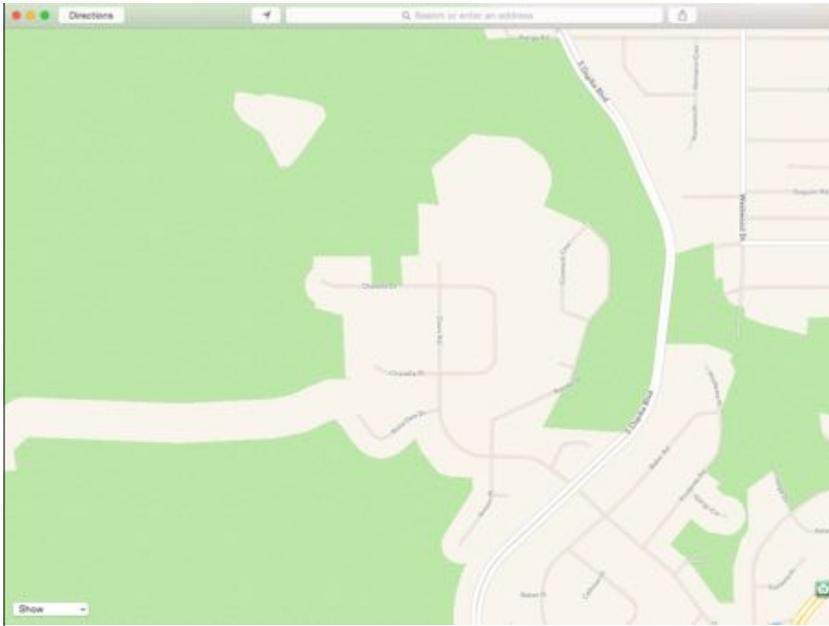
“It’s hobbyist  
stuff”

between the Encyclopedia Britannica and Wikipedia, one has always had the lock on "authoritative" but the other one is what **we actually use to look up an answer.**

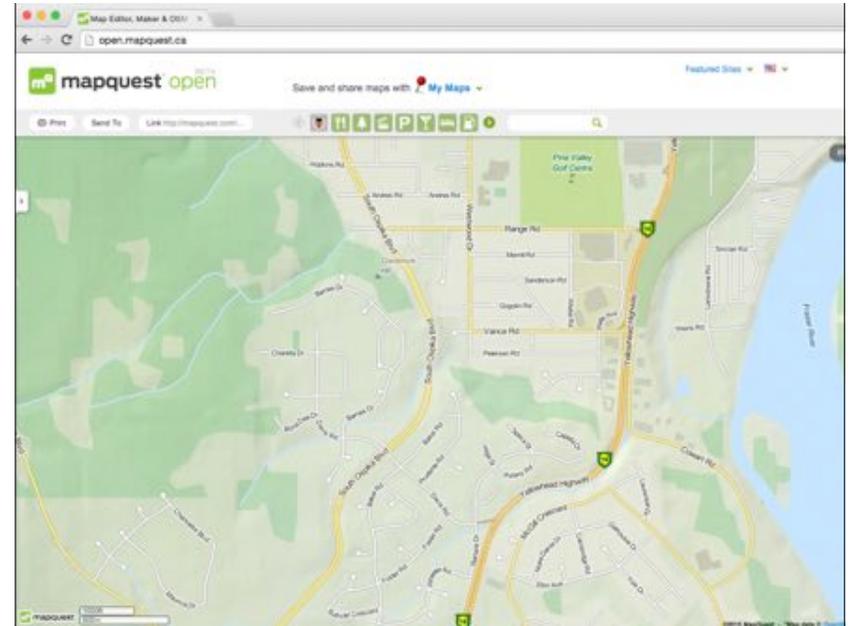
So being "authoritative" is mostly meaningless in the relevance sweepstakes, and hopefully everyone here is interested in remaining relevant.

"it's hobbyist stuff"

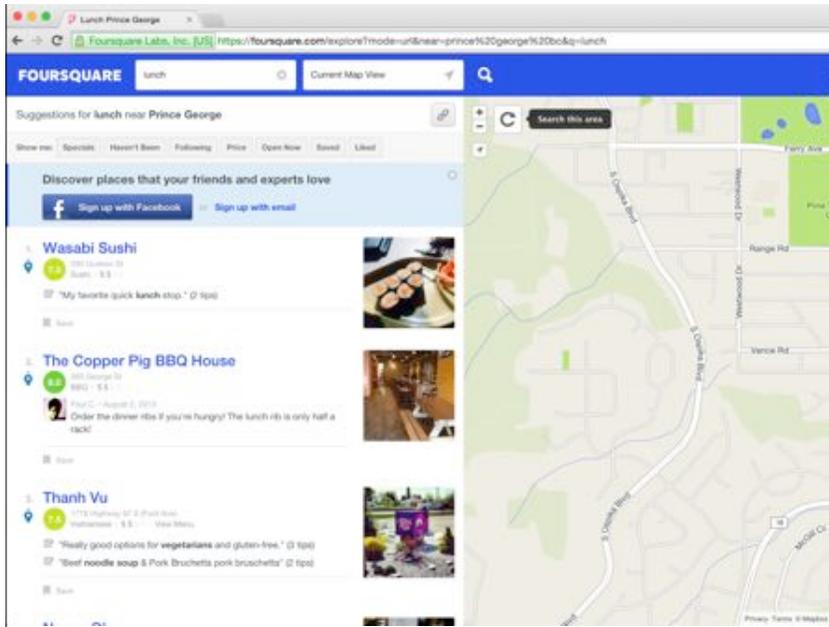
On hobbyist, those days are long, long past.



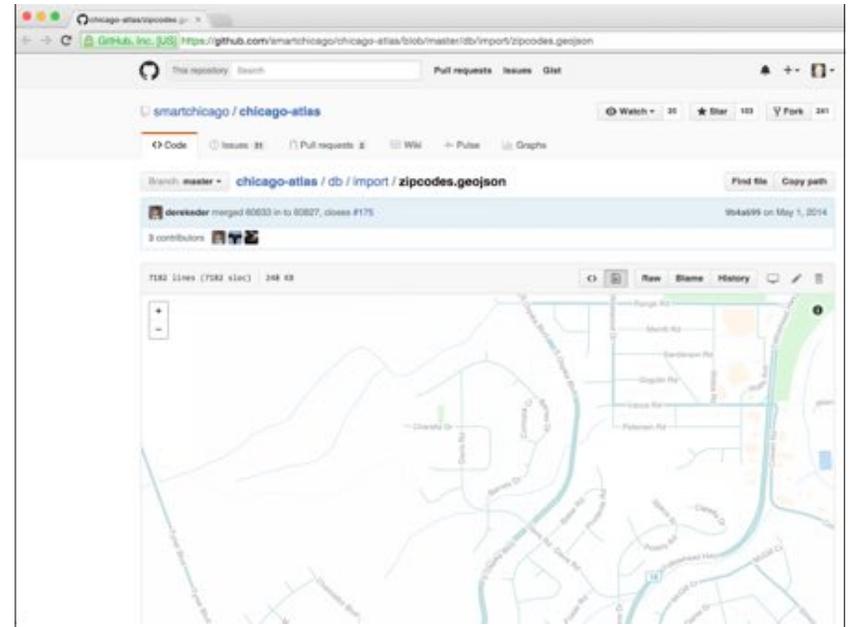
- Apple maps makes partial use of openstreetmap data
- Tableau provides OpenStreetMap map tiles they render themselves



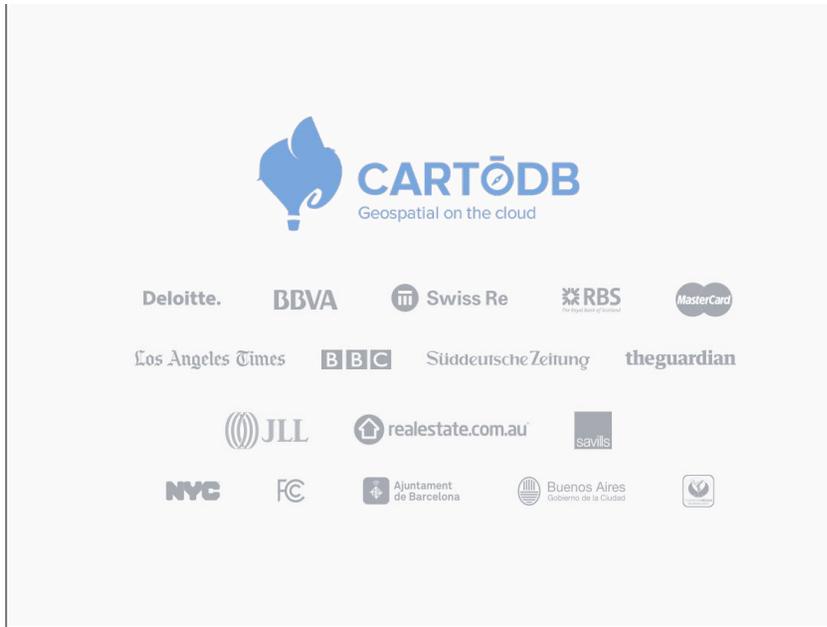
- MapQuest makes extensive use of openstreetmap



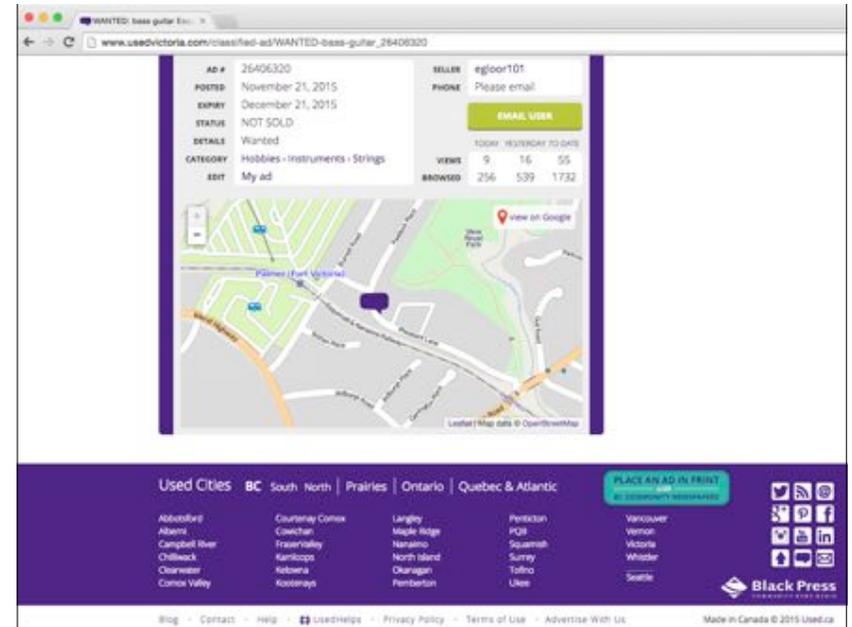
- Foursquare and Pinterest and National Geographic, Road Tripper, the Financial Times, The Guardian, USA Today, the Wall Street Journal,



gitHub, Etsy, the FCC and the Washington Post are only \*some\* of the institutional and corporate customers who use OpenStreet Map on their public maps, via map tiles produced by MapBox



- Everyone of our 200,000 users at CartoDB, including City of New York, NewsWeek, the LA Times, UN Environmental Program, and the US National Parks Service, makes use of Openstreetmap in the basemaps we provide for free with our data visualization and analysis services



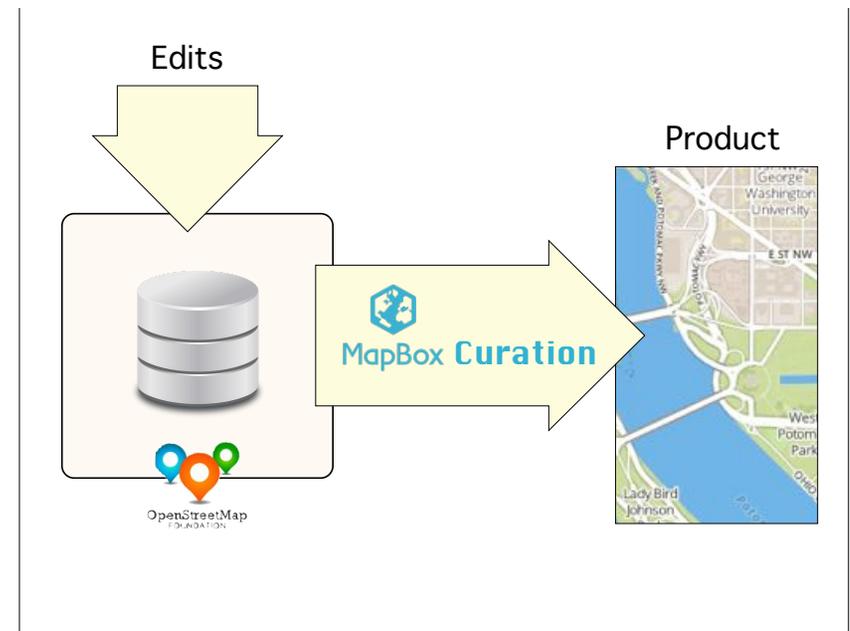
- I couldn't help noticing that here in Victoria even our own UsedVictoria.com uses OSM tiles (it looks like they use the community rendering, so it's kind of obvious)

“It changes!”  
“It’s wrong!”

"it changes!" / "it's wrong!"

Sure, OSM changes, sometimes very fast, and sometimes it's wrong.

Mostly it's right though (studies have been done), and it's not hard to make use of it with mitigating strategies to deal with the velocity of change.



MapBox manages the flow on behalf of its clients by reviewing the edit stream before applying it to their publishing database. That requires a lot of human labour, but a lot less labour than required to generate the edit stream in the first place.

The question isn't if OSM changes or if OSM is wrong,

“Whaddaya gonna do about it?”

The question is “Whaddaya gonna do about it?”

Because OSM is already in place and sucking up all the oxygen.

You can figure out how to collaborate and work with the OSM ecosystem,



or you stand still and shout at it as your relevance slowly vanishes.

It's worth noting that a few years ago Britannica had to **exit** the encyclopedia business.

They stopping printing encyclopedias, you can no longer buy a paper copy of Britannica.

Authoritative



Relevant



WIKIPEDIA  
The Free Encyclopedia

# Geodetic Control

Not because their encyclopedias had gotten any worse over the years,  
not because their data was no good,  
Just because  
nobody wanted  
to use them  
anymore.

If I had more time,  
I would talk about the threats to the relevance of  
government work in geodetic control and  
registration of data,  
but then I wouldn't be able to  
complete the thought about relevance in  
distribution.

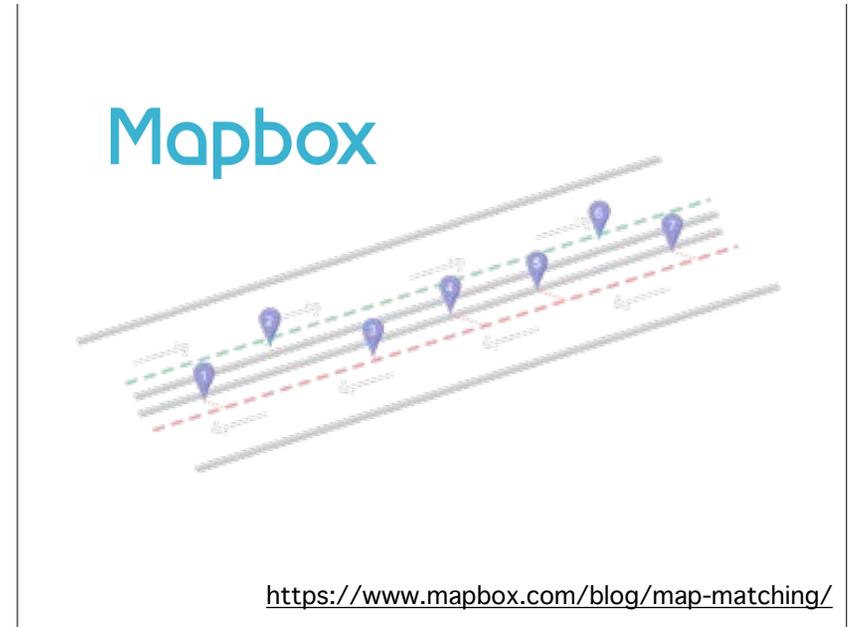


Suffice it to say,  
the start-up world is increasingly bypassing explicit sources of control  
and using machine learning technology to register new data to  
stacks of old data.

I cannot recommend enough a talk from PlanetLabs  
about how they are registering and rectifying  
over 1MILLION scenes a day with no human input at all.

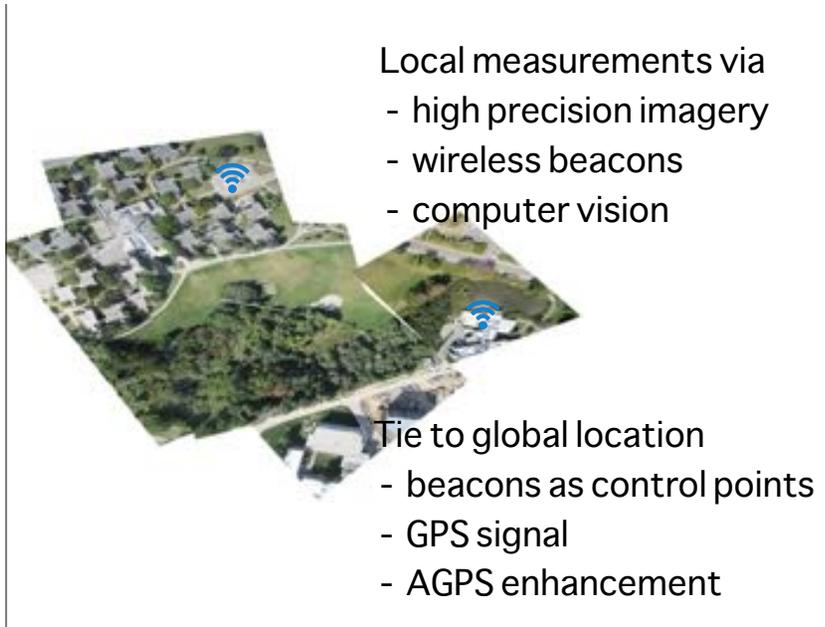
It's called rectifying the planet, and it's on the reading list page.

<https://2015.foss4g-na.org/session/rectifying-planet>



<https://www.mapbox.com/blog/map-matching/>

Similarly, a blog post from mapbox about  
statistical conflation of GPS track data is well worth  
reading.  
It's also on the reading list



The ability to build accurate models from image collections means that relative measurement can be done on the fly, extremely accurately without any special data gathering effort.

So high precision local measurement is getting better all the time. And wireless beacons, like wifi signals are increasingly being used to tie local data into a global frame in real time.

Over time, all this new technology is going to bypass traditional flows for tying data into a global framework, which may erode the relevance of local geodetic frameworks.



Which brings me to cadastre,  
the last bastion of government control.  
And I don't think I'm overstating things much  
to say government is hard at work screwing this  
up,  
setting the stage for a final loss of relevance.

## Cadastre is valuable: why?

---

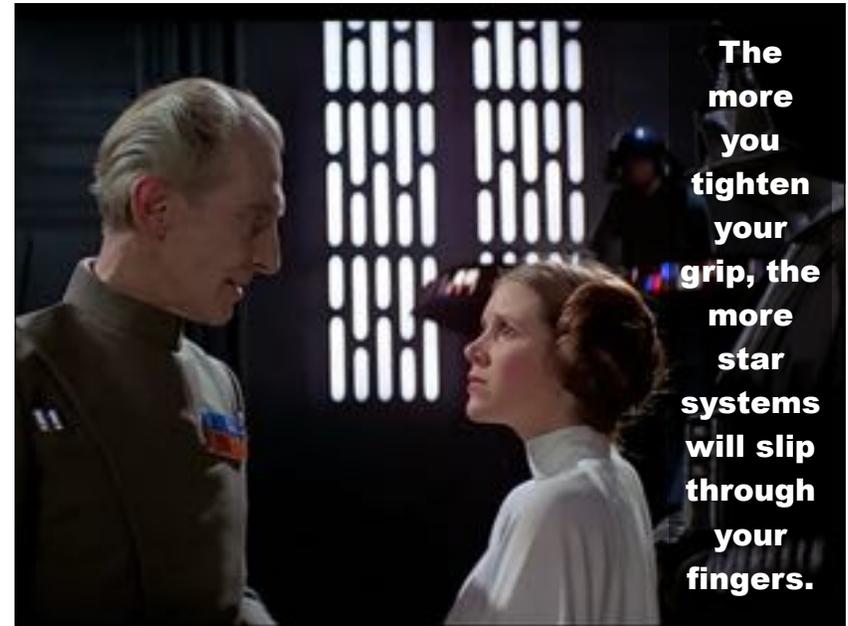
scarcity

control

validity

Cadastre has "value", a value predicated on scarcity and control and volatility. And governments continue to use that value to **leverage revenue** from the data, and protect certain classes of businesses and professionals through preferential access to the data.

This won't last forever.



Princess Leia was right. The more you tighten your grip, the more star system will slip through your fingers.

It's worth remembering where OSM came from, this mapping project that has basically taken over the physical mapping space.



It came about because the UK Ordnance Survey maintained a **monopoly** on physical mapping in the UK, and leveraged as **much revenue as possible** from that monopoly, charging what municipalities and regions and corporations were able to pay for the data.

They thought they were doing OK, because their "clients" were happy.

The only unhappy people weren't clients, they were whiney academics and civil society types who couldn't afford the data, or wanted to process it in ways the Ordnance Survey license wouldn't allow.



And eventually, not because they necessarily thought they would succeed, but more to give the finger to the Ordnance survey, a couple students spent a month building a basic infrastructure for crowd sourcing maps, and then started telling people about how to do it.

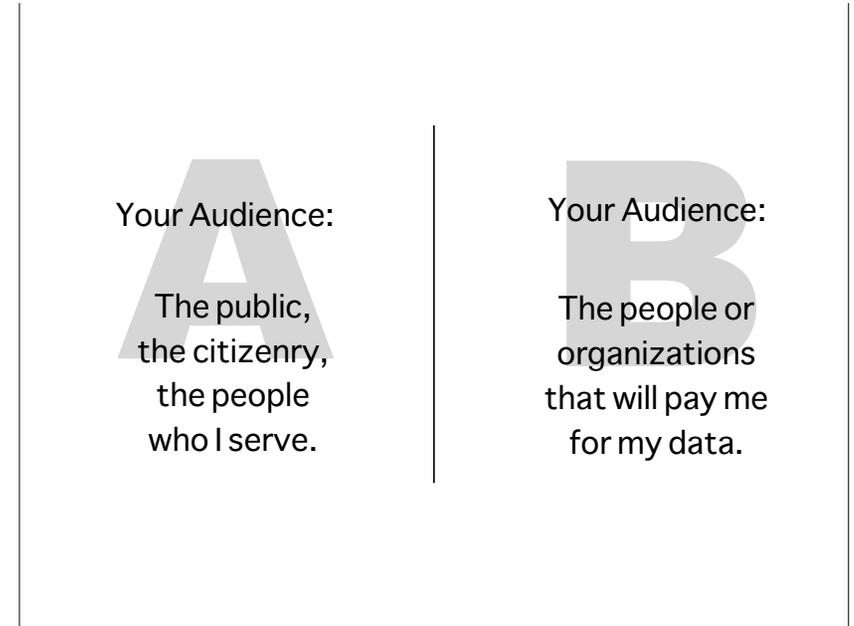
Mapping their local neighborhoods, building an unencumbered version of the data that anyone could use, any way they wanted to.



And that's what they did, going from nothing, to a crowd-sourced map of the world that is now a one-stop shop to anyone building a location application.

It was exceedingly improbable when it started, and yet, here it is now.

Still cadastre looks like an impregnable bastion of control right now, if only because of the legal structures that bring it into existence and the general lack of external observability of legal boundaries.



That shouldn't be the end of the story though, because the purpose of government, is to serve the public, by providing public goods. At least I think so. And I think the public should be your primary audience.

But maybe you define your audience differently <X> If you are defining the audience as "those people who will pay for my cadastre", you may be making a huge category error, an error that will erode your relevance in the long term, and set the stage for some future crowd-sourced project to take you down.



Is what you build  
getting used?

The people who built this freeway may have been very proud of it. Look at all the lanes, look at the quality of the paving. But there's nobody on it. Maybe because there's nobody around, maybe because the tolls are too high.

If the public isn't using a public good, it's a misapplied expenditure. It's misapplied effort to build data that nobody gets to use, even really excellent data, of the highest quality and currency.

And that's where much government cadastre sits right now, either because of old cost recovery promises, or because of institutional restrictions invented to protect professional societies.

## Distributing Cadastre

But supposing you can break out of the current data black hole enveloping cadastre, what distribution scheme makes sense?

A completely community oriented model like OSM probably doesn't fit for distribution of cadastre, because it's very difficult to find an alternate source of truth for cadastre:  
the government knows where parcels are, and that's about it.

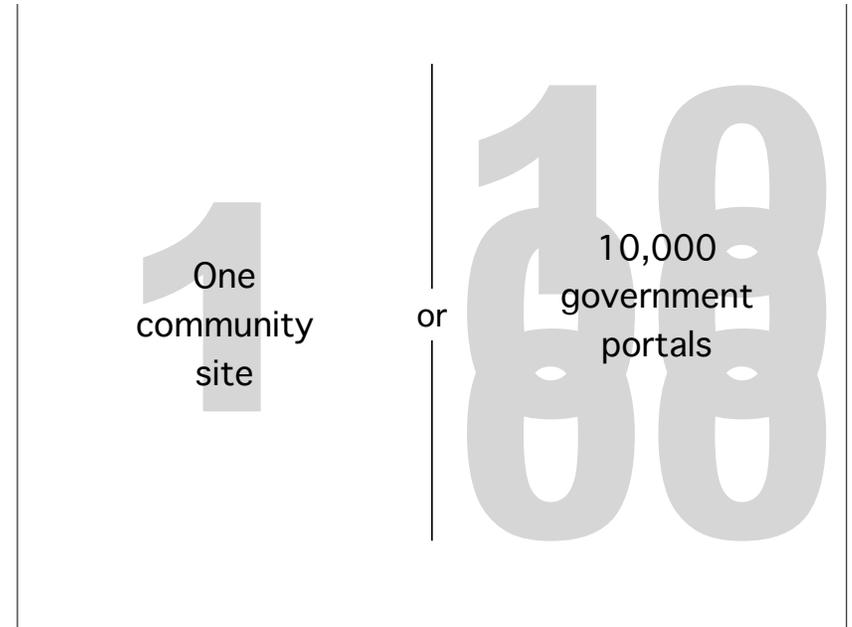
However, there's already a project in place that demonstrates a different, more appropriate model, and as a bonus, it's a really cool approach.



OpenAddresses is a community project that doesn't manage raw data directly. Rather, it manages metadata necessary to assemble raw data into a consistent data set.

A contributor to open-addresses doesn't have to contribute an address file but can instead contribute a **recipe** to download the data from the agency that created it, and normalize that data into the standard openaddresses schema.

The result is a recipe box that can be run to create the openaddresses corpus, which currently includes over 218 million addresses.



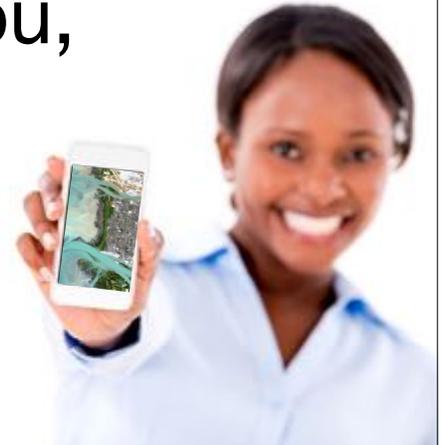
It's easy to imagine a companion project that collated cadastre information, using the same principles.

And It's easy to imagine the public needs such a project would meet, at the fraction of a cost of standing up data portals and fancy API systems.

# Moment of Opportunity

This is a moment of opportunity, as well a moment of threat.

If not you,  
who?



If you don't put government data on those devices, someone else will put something else there instead, and your citizens, your elected officials, will be wondering, what do you folks do, anyways?

So, what's the road to get onto the devices, to get in front of those eyeballs?

$$y = x \cdot \text{change}$$

Developers aren't looking directly to government anymore, but that doesn't mean the game is up, it just means that governments need to accept the way that the technology ecosystem is going to want to consume their data, and change their behavior to fit.

## Data as Public Good

The first step, is to recommit to the idea of data as a public good. If this data is critical infrastructure, as we believe it to be, making it available to all members of civil society, without restriction, is a basic requirement.

benefit ~ reach

Once that is done, strategic directions become a lot clearer, because the question becomes one of how to maximize the reach of the data.

How to get data into the pipes that serve all those mobile devices, and all the people carrying them about in their pockets.

And that's a very solvable problem.

**! Portals  
! APIS**

(OK, nerd nomenclature break. Using an ! as a prefix is a negation symbol in most programming languages.) Soo...

Can we all agree to take a break for a while and stop standing up data clipping portals and database APIs?

We have enough bandwidth to ship and process data

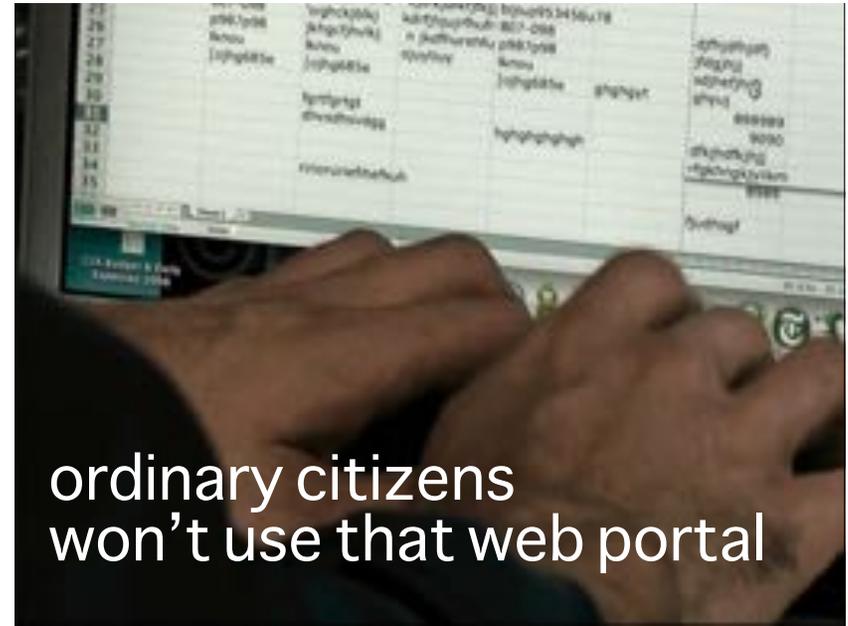
in multi-gigabyte sized chunks, so we don't need to clip and ship anymore. Stop.

# serve your data



## raw

Commit to simplicity in distribution. Follow the lead of NASA and publish raw data, with computer readable manifests, with stable URLs, close to the point of consumption on public cloud infrastructure.



## ordinary citizens won't use that web portal

Too much effort has been wasted trying to make open data usable by "ordinary citizens". Most open data will only be accessed by specialists, so let's accept that, and publish bulk data first, not summaries, and not charting tools.

Open data catalogues remain useful in providing a single search location for data, but there's no reason the raw data cannot exist in the cloud, with metadata records pointing to the appropriate documentation on the storage scheme.

# Engage Existing Communities

This one will be the hardest to accept and to implement, but it's critical to avoiding irrelevance in the long term as mapping communities continue to grow.

A common repository  
of all physical mapping  
in the world ...

- A common repository of all physical mapping in the world ...



already exists, it's called open street map  
- If you put your data in there, it will find its way to your citizens, guaranteed



- Because the pipes of data usage running out from OSM are already huge and growing  
There's no better route to getting onto people's devices than using OSM as a distribution channel.  
None.  
And I know it will be hard. There will be excuses.

“We can't use OSM.  
Because: the law”

"We can't use OSM. Because: the law".  
Speaking as someone standing on the outside looking in,  
there's few things as galling as the people who make the laws  
using the current state of the laws as an excuse for inaction.

I can accept that the current legal frameworks make things **difficult**,  
but not without pointing out



that the only way to guarantee that things won't change is to **not** try to change them.  
(That's the Little Engine that Could, by the way.)

And if you **do not try** make the change,  
you are accepting the current trajectory,  
of increasing irrelevance over time.

When the facts change, it's time to change your mind.  
The facts of spatial data have changed.

# Build New Communities

If you accept (as I hope you might) that distribution via cooperative infrastructures, like OSM or OpenAddresses, is the wave of the future, you'll note that there are some gaps in the current ecosystem.



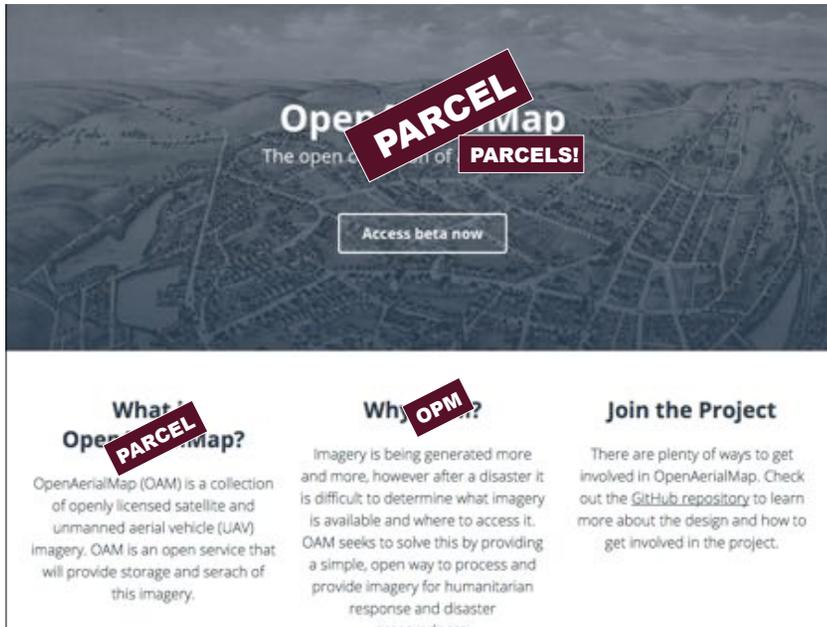
There is a community for distribution of imagery, OpenAerialMap but it is still very much a prototype, with limited infrastructure and participation.

A shared repository of the world's free imagery would be an immense benefit to the citizens of every nation, but the effort has practically no resources,

it's been bootstrapped by the "humanitarian openstreetmap team", with a small amount of World Bank funding.

This is an opportunity for Canada to show world leadership, and help our own data distribution needs at the same time.

It's really a crime that such an important effort is left to volunteers and nonprofits.



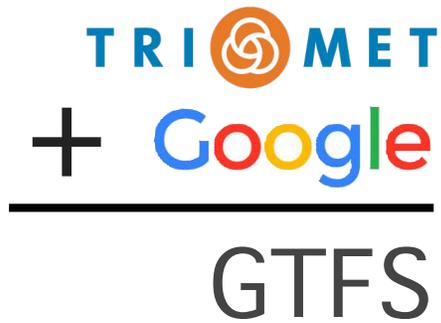
As mentioned, there's no community solution at all for sharing cadastre data.  
We can wait for one to show up organically,  
or we can try to shape one that is most pleasing to us,  
perhaps using the OpenAddresses model of integration recipes.

This is another place where relatively few resources could be applied with immense leverage.

<http://wiki.openstreetmap.org/wiki/Parcel>

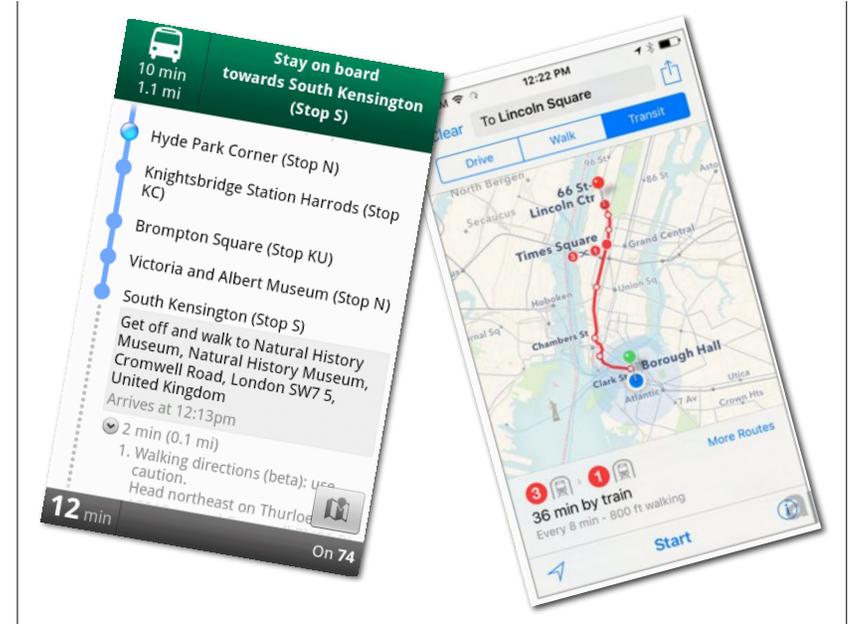


It's possible for a public agency to have a huge positive world-wide effect,  
just through the power of promoting an open solution.



[https://en.wikipedia.org/wiki/General\\_Transit\\_Feed\\_Specification](https://en.wikipedia.org/wiki/General_Transit_Feed_Specification)

When Portland Trimet was looking to publish their transit schedule data for the first time, they didn't just stick it on the web, or buy some COTS solution and call it a day, <X> they looked for a big potential consumer of the data, Google, and asked them what the most useful open format would be. <X> The result was the General Transit Feed Specification (GTFS). When you run a transit routing search on Google Maps now,



it works uniformly across hundreds of transit systems, because of GTFS. And the availability of open GTFS format data means transit routing is available on multiple platforms and apps, not just Google's.



The format has spawned numerous third-party uses, from "next bus" LED signs, to apps, to open source routing solutions like OpenTripPlanner, and education materials to learn more.

One transit agency seeded the whole thing, by committing to openness, to getting their data in front of the maximum number of eyeballs and recognizing **their needs were not unique.**



Government mapping shouldn't be declining in perceived relevance, but it is. You can reverse it, though. Give up some direct control, in exchange for wider distribution, getting in front of more eyeballs. Getting in front of your citizens, your politicians, the people whose opinions about your relevance really matter. Engage with the world.

# Thank You

<http://j.mp/ccog2015>

# Sorry