

EXHIBIT 3

The meaning of open

Monday, December 21, 2009 at 4:57 PM ET

Posted by Jonathan Rosenberg, Senior Vice President, Product Management

Last week I sent an email to Googlers about the meaning of "open" as it relates to the Internet, Google, and our users. In the spirit of openness, I thought it would be appropriate to share these thoughts with those outside of Google as well.

At Google we believe that open systems win. They lead to more innovation, value, and freedom of choice for consumers, and a vibrant, profitable, and competitive ecosystem for businesses. Many companies will claim roughly the same thing since they know that declaring themselves to be open is both good for their brand and completely without risk. After all, in our industry there is no clear definition of what open really means. It is a Rashomon-like term: highly subjective and vitally important.

The topic of open seems to be coming up a lot lately at Google. I've been in meetings where we're discussing a product and someone says something to the effect that we should be more open. Then a debate ensues which reveals that even though most everyone in the room believes in open we don't necessarily agree on what it means in practice.

This is happening often enough for me to conclude that we need to lay out our definition of open in clear terms that we can all understand and support. What follows is that definition based on my experiences at Google and the input of several colleagues. We run the company and make our product decisions based on these principles, so I encourage you to carefully read, review, and debate them. Then own them and try to incorporate them into your work. This is a complex subject and if there is debate (and I'm sure there will be) it should be in the open! Please feel free to comment.

There are two components to our definition of open: open technology and open information. Open technology includes open source, meaning we release and actively support code that helps grow the internet, and open standards, meaning we adhere to accepted standards and, if none exist, work to create standards that improve the entire internet (and not just benefit Google). Open information means that when we have information about users we use it to provide something that is valuable to them, we are transparent about what information we have about them, and we give them ultimate control over their information. These are the things we should be doing. In many cases we aren't there, but I hope that with this note we can start working to close the gap between reality and aspiration.

If we can embody a consistent commitment to open --- which I believe we can --- then we have a big opportunity to lead by example and encourage other companies and industries to adopt the same commitment. If they do, the world will be a better place.

Open systems win

To understand our position in more detail, it helps to start with the assertion that open systems win. This is counter-intuitive to the traditionally trained MBA who is taught to generate a sustainable competitive advantage by creating a closed system, making it popular, then milking it through the product life cycle. The conventional wisdom goes that companies

...blogspot.com/.../meaning-of-ope...

Search This Blog

powered by Google™

Site Feed

Analytics

7994 readers

More Blogs from Google

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

Get Blog Posts by E-mail

[Intro to Google for Policymakers](#)

Follow us on Twitter



Google Public Policy
googlepubpolicy

Google exec Jonathan Rosenberg on the meaning of open: <http://bit.ly/4e2i6X>
3 hours ago

From the blog: Google's approach to privacy: <http://bit.ly/30cKld>
3 days ago

twitter

Join the conversation

Archive

What We're Reading

[463 Blog: Inside Tech Policy](#)

[CDT - PolicyBeta](#)

12/22/2009

Google Public Policy Blog: The me...

should lock in customers to lock out competitors. There are different tactical approaches — razor companies make the razor cheap and the blades expensive, while the old IBM made the mainframes expensive and the software ... expensive too. Either way, a well-managed closed system can deliver plenty of profits. They can also deliver well-designed products in the short run — the iPod and iPhone being the obvious examples — but eventually innovation in a closed system tends towards being incremental at best (is a four blade razor really that much better than a three blade one?) because the whole point is to preserve the status quo. Complacency is the hallmark of any closed system. If you don't have to work that hard to keep your customers, you won't.

Open systems are just the opposite. They are competitive and far more dynamic. In an open system, a competitive advantage doesn't derive from locking in customers, but rather from understanding the fast-moving system better than anyone else and using that knowledge to generate better, more innovative products. The successful company in an open system is both a fast innovator and a thought leader; the brand value of thought leadership attracts customers and then fast innovation keeps them. This isn't easy — far from it — but fast companies have nothing to fear, and when they are successful they can generate great shareholder value.

Open systems have the potential to spawn industries. They harness the intellect of the general population and spur businesses to compete, innovate, and win based on the merits of their products and not just the brilliance of their business tactics. The race to map the human genome is one example.

In the book *Wikinomics*, Don Tapscott and Anthony Williams explain how in the mid-1990s private firms were discovering and patenting large amounts of DNA sequence data and then assuming control over who could access that information and at what price. Having so much of the genome under private ownership raised costs and made drug discovery far less efficient. Then, in 1995, Merck Pharmaceuticals and the Gene Sequencing Center at Washington University changed the game by creating a new, open initiative called the Merck Gene Index. Within three years they had published over 800,000 gene sequences into the public domain, and soon other collaborative projects followed suit. This in an industry where early stage R&D was traditionally pursued individually in closed labs, so Merck's open approach not only changed the culture of the entire field but also accelerated the pace of biomedical research and drug development. It gave researchers everywhere unrestricted access to an open resource of genetic information.

Another way to look at the difference between open and closed systems is that open systems allow innovation at all levels — from the operating system to the application layer — not just at the top. This means that one company doesn't have to depend on another's benevolence to ship a product. If the [GNU C compiler](#) that I'm using has a bug, I can fix it since the compiler is open source. I don't have to file a bug report and hope for a timely response.

So if you are trying to grow an entire industry as broadly as possible, open systems trump closed. And that is exactly what we are trying to do with the Internet. Our commitment to open systems is not altruistic. Rather it's good business, since an open Internet creates a steady stream of innovations that attracts users and usage and grows the entire industry. Hal Varian has an equation in his book *Information Rules* that applies here:

$$\text{Reward} = (\text{Total value added to the industry}) * (\text{Our share of industry value})$$

All other things being equal, a 10 percent increase in share or a 10 percent increase in industry value should lead to the same outcome. But in our industry a 10 percent increase in industry value will yield a much bigger reward because it will stimulate economies of scale across the entire industry, increasing productivity and reducing costs for all competitors. As long as we contribute a steady stream of great products we will prosper along with the entire ecosystem. We may get a smaller piece, but it will come from a bigger pie.

In other words, Google's future depends on the Internet staying an open system, and our advocacy of open will grow the web for everyone - including Google.

Open Technology

The definition of open starts with the technologies upon which the Internet was founded: open standards and open source software.

...blogspot.com/.../meaning-of-ope...

[Cisco High Tech Policy Blog](#)
[CNET Politics Blog](#)
[EFF DeepLinks](#)
[GigaOM](#)
[Google European Policy Blog](#)
[Google Public Sector Blog](#)
[Hillicon Valley](#)
[Intel Policy Blog](#)
[John Palfrey's Blog](#)
[Lessig Blog](#)
[Microsoft On the Issues](#)
[Peter Fleischer: Privacy...?](#)
[PFF Blog](#)
[Post I.T.](#)
[Public Knowledge Blog](#)
[Save the Internet](#)
[Search Engine Land](#)
[Susan Crawford Blog](#)
[Tech Daily Dose](#)
[Tech Dirt](#)
[TechCrunch](#)
[Technology Liberation Front](#)
[Verizon Policy Blog](#)
[Yahoo! Policy Blog](#)

Labels

[Accessibility](#) (1)
[Advertising](#) (5)
[Book Search](#) (15)
[Business Issues](#) (22)
[Canada](#) (1)
[Child Safety](#) (13)
[Competition](#) (10)
[Cybersecurity](#) (4)
[D.C. Talks](#) (12)
[Economy](#) (10)
[Elections](#) (23)
[Energy Efficiency](#) (22)
[Free Expression](#) (25)
[Google Tools](#) (55)
[Government Transparency](#) (13)
[Health](#) (4)
[Intellectual Property](#) (14)
[International](#) (43)
[Net Neutrality](#) (19)
[Politicians at Google](#) (11)
[Politics](#) (19)

2/9

12/22/2009

Google Public Policy Blog: The me...

Open Standards

Networks have always depended on standards to flourish. When railroad tracks were first being laid across the U.S. in the early 19th century, there were seven different standards for track width. The network didn't flourish and expand west until the different railway companies agreed upon a standard width of 4' 8.5". (In this case the standards war was an actual war: Southern railroads were forced to convert over 11,000 miles of track to the new standard after the Confederacy lost to the Union in the Civil War.)

So there was some precedent in 1974 when Vint Cerf and his colleagues proposed using an open standard (which became TCP/IP) to connect the several computer networks that had emerged around the U.S. They didn't know exactly how many networks were out there so the "Internet" — a term Vint coined — had to be open. Any network could connect using TCP/IP, and now, as a result of that decision, there are about 681 million hosts on the Internet.

Today, we base our developer products on open standards because interoperability is a critical element of user choice. What does this mean for Google Product Managers and Engineers? Simple: whenever possible, use existing open standards. If you are venturing into an area where open standards don't exist, create them. If existing standards aren't as good as they should be, work to improve them and make those improvements as simple and well documented as you can. Our top priorities should always be users and the industry at large and not just the good of Google, and you should work with standards committees to make our changes part of the accepted specification.

We have a good history of doing this. In the formative years of the [Google Data Protocol](#) (our standard API protocol, which is based on XML/Atom), we worked as part of the IETF Atom Protocol Working Group to shape the Atom specification. There's also our recent work with the W3C to create a standard geolocation API that will make it easy for developers to build browser-based, location-sensitive applications. This standard helps everyone, not just us, and will lead to users having access to many more compelling apps from thousands of developers.

Open Source

Most of those apps will be built on open source software, a phenomenon responsible for the web's explosive growth in the past 15 years. There is a historic precedent here: while the term "open source" was coined in the late 1990s, the concept of sharing valuable information to catalyze an industry existed long before the Internet. In the early 1900s, the U.S. automobile industry instituted a cross-licensing agreement whereby patents were shared openly and freely amongst manufacturers. Prior to this agreement, the owners of the patent for the two-cycle gasoline engine had effectively bottled up the industry.

Today's open source goes far beyond the "patent pooling" of the early auto manufacturers, and has led to the development of the sophisticated software components — Linux, Apache, SSH, and others — upon which Google is built. In fact, we use tens of millions of lines of open source code to run our products. We also give back: we are the largest open source contributor in the world, contributing over 800 projects that total over 20 million lines of code to open source, with four projects (Chrome, Android, Chrome OS, and Google Web Toolkit) of over a million lines of code each. We have teams that work to support Mozilla and Apache, and an open source project hosting service (code.google.com/hosting) that hosts over 250,000 projects. These activities not only ensure that others can help us build the best products, they also mean that others can use our software as a base for their own products if we fail to innovate adequately.

When we open source our code we use standard, open [Apache 2.0 licensing](#), which means we don't control the code. Others can take our open source code, modify it, close it up and ship it as their own. [Android](#) is a classic example of this, as several OEMs have already taken the code and done great things with it. There are risks to this approach, however, as the software can fragment into different branches which don't work well together (remember how Unix for workstations devolved into various flavors — Apollo, Sun, HP, etc.). This is something we are working hard to avoid with Android.

While we are committed to opening the code for our developer tools, not all Google products are open source. Our goal is to keep the Internet open, which promotes choice and competition and keeps users and developers from getting locked in. In many cases, most notably our search and ads products, opening up the code would not contribute to these goals and would actually hurt users. The search and advertising markets are already highly

...blogspot.com/.../meaning-of-ope...

[Privacy](#) (52)

[Security](#) (1)

[State Issues](#) (5)

[Telecom](#) (50)

[White Spaces](#) (16)

[Workforce](#) (4)

[Yahoo-Google Deal](#) (5)

This blog is powered by Blogger.

[Start your own weblog.](#)



12/22/2009

Google Public Policy Blog: The me...

competitive with very low switching costs, so users and advertisers already have plenty of choice and are not locked in. Not to mention the fact that opening up these systems would allow people to "game" our algorithms to manipulate search and ads quality rankings, reducing our quality for everyone.

So as you are building your product or adding new features, stop and ask yourself: Would open sourcing this code promote the open internet? Would it spur greater user, advertiser, and partner choice? Would it lead to greater competition and innovation? If so, then you should make it open source. And when you do, do it right; don't just push it over the wall into the public realm and forget about it. Make sure you have the resources to pay attention to the code and foster developer engagement. [Google Web Toolkit](#), where we have developed in the open and used a public bug tracker and source control system, is a good example of this.

Open Information

The foundation of open standards and open source has led to a web where massive amounts of personal information — photos, contacts, updates — are regularly uploaded. The scale of information being shared, and the fact that it can be saved forever, creates a question that was hardly a consideration a few years ago: How do we treat this information?

Historically, new information technologies have often enabled new forms of commerce. For example, when traders in the Mediterranean region circa 3000 BC invented seals (called bullae) to ensure that their shipments reached their destinations tamper-free, they transformed commerce from local to long distance. Similar transformations were spurred by the advent of the written word, and more recently, computers. At every step of the way, the transaction, a consensual agreement where each party gets something of value, was powered by a new type of information that allowed a contract to be enforced.

On the web, the new form of commerce is the exchange of personal information for something of value. This is a transaction that millions of us participate in every day, and it has potentially great benefits. An auto insurer could monitor a customer's driving habits in real-time and give a discount for good driving — or charge a premium for speeding — powered by information (GPS tracking) that wasn't available only a few years ago. This is a fairly simple transaction, but we will encounter far more sensitive scenarios.

Let's say your child has an allergy to certain medicines. Would you allow her medical data to be accessible by a smart wireless syringe which could prevent an EMT or nurse from accidentally giving her that medicine? I would, but you might decide the metal bracelet around her wrist is sufficient. And that's the point — people can and will reach different decisions, and when it comes to their personal information we need to treat all of those decisions with equal respect.

So while having more personal information online can be quite beneficial to everyone, its uses should be guided by principles that are responsible, scalable, and flexible enough to grow and change with our industry. And unlike open technology, where our objective is to grow the internet ecosystem, our approach to open information is to build trust with the individuals who engage within that ecosystem (users, partners, and customers). Trust is the most important currency online, so to build it we adhere to three principles of open information: value, transparency, and control.

Value

First and foremost, we need to make products that are valuable to users. In many cases, we can make our products even better if we know more information about the user, but privacy concerns can arise if people don't understand what value they are getting in return for their information. Explain that value to them, however, and they will often agree to the transaction. For example, millions of people let credit card companies retain information on the purchases they make with their card in exchange for the convenience of not carrying around cash.

We did this well when we launched [Interest-Based Advertising](#) in March. IBA makes ads more relevant and more useful. That is the extra value we create based on the information we gather. It also includes a [user preferences manager](#) that clearly explains what users are getting in exchange for their information and lets them opt out or adjust their settings. The vast majority of people who visit the preferences manager choose to adjust their settings rather than opt out because they realize the value of receiving ads customized to their interests.

...blogspot.com/.../meaning-of-ope...

4/9

12/22/2009

Google Public Policy Blog: The me...

This should be our default approach: tell people, in obvious, plain language, what we know about them and why it's valuable to them that we know it. Think that your product's value is so obvious that it doesn't need explaining? There's a good chance you're wrong.

Transparency

Next, we need to make it easy for users to find out what information we gather and store about them across all of our products. We recently took a big step in this direction with the launch of the Google Dashboard, which is a single place where users can see what personal data is held by each Google product (covering more than 20 products including Gmail, YouTube, and Search) and control their personal settings. We are, to the best of our knowledge, the first internet company to offer a service like this and we hope it will become the standard. Another good example is our [Privacy Policy](#), which is written for humans and not just lawyers.

We can go even farther than this though. If you manage a consumer product where you collect information from your users, your product should be part of the Dashboard. If you're already there, you're not done. With every new feature or version, ask yourself if you have any additional information (maybe even information that is publicly available about users on other sites) that you can add to the Dashboard.

Think about how you can increase transparency within your product as well. When you download an Android app, for example, the device tells you what information the app will be able to access about you and your phone, and then you get to decide whether or not to proceed. You don't have to dig deep to figure out what information you are divulging - it tells you up front and lets you decide what to do. Is your product like that? How can you increase users' engagement with your product through increasing transparency?

Control

Finally, we must always give control to the user. If we have information about a user, as with iBA, it should be easy for the user to delete that information and opt-out. If they use our products and store content with us, it's their content, not ours. They should be able to export it or delete it at any time, at no cost, and as easily as possible. Gmail is a great example of this since we offer free forwarding to any address. The ability to switch is critical, so instead of building walls around your product, build bridges. Give users real options.

If there are existing standards for handling user data, then we should adhere to them. If a standard doesn't exist, we should work to create an open one that benefits the entire web, even if a closed standard appears to be better for us (remember — it's not!). In the meantime we need to do whatever we can to make leaving Google as easy as possible. Google is not the Hotel California — you can check out any time you like and you CAN, in fact, leave!

As Eric said in his 2009 strategy memo, "we don't trap users, we make it easy for them to move to our competitors." This policy is sort of like the emergency exits on an airplane — an analogy that our pilot CEO would appreciate. You hope to never use them, but you're glad they're there and would be furious if they weren't.

That's why we have a team — the Data Liberation Front (dataliberation.org) — whose job it is to make "checking out" easy. Recent examples of their work include Blogger (people who choose to leave Blogger for another service can easily take their content with them) and Docs (users can now collect all their documents, prescos, and spreadsheets in a zip file and download it). Build your products so that the Data Liberation team can work their magic. One way you can do this is by having a good public API that exposes all your users' data. Don't wait for v2 or v3, discuss this early in your product planning meetings and make it a feature of your product from the start.

When reporters at the Guardian, a leading UK newspaper, [reviewed the work](#) of the Data Liberation team they proclaimed it to be "counter-intuitive" for those "accustomed to the lock-in mentality of previous commercial battles." They are right, it is counterintuitive to people who are stuck in the old MBA way of thinking, but if we do our jobs then soon it won't be. Our goal is to make open the default. People will gravitate towards it, then they will expect and demand it and be furious when they don't get it. When open is intuitive, then we have succeeded.

When bigger is better

Closed systems are well-defined and profitable, but only for those who control them. Open systems are chaotic and profitable, but only for those who understand them well and move

...blogspot.com/.../meaning-of-ope...

5/9

12/22/2009

Google Public Policy Blog: The me...

faster than everyone else. Closed systems grow quickly while open systems evolve more slowly, so placing your bets on open requires the optimism, will, and means to think long term. Fortunately, at Google we have all three of these.

Because of our reach, technical know-how, and lust for big projects, we can take on big challenges that require large investments and lack an obvious, near-term pay-off. We can photograph the world's streets so that you can explore the neighborhood around an apartment you are considering renting from a thousand miles away. We can scan millions of books and make them widely accessible (while respecting the rights of publishers and authors). We can create an email system that gives away a gigabyte of storage (now over 7 gigs) at a time when all other services allow only a small fraction of that amount. We can instantly translate web pages from any of 51 languages. We can process search data to help public health agencies detect flu outbreaks much earlier. We can build a faster browser (Chrome), a better mobile operating system (Android), and an entirely new communications platform (Wave), and then open them up for the world to build upon, customize, and improve.

We can do these things because they are information problems and we have the computer scientists, technology, and computational power to solve them. When we do, we make numerous platforms - video, maps, mobile, PCs, voice, enterprise - better, more competitive, and more innovative. We are often attacked for being too big, but sometimes being bigger allows us to take on the impossible.

All of this is useless, however, if we fail when it comes to being open. So we need to constantly push ourselves. Are we contributing to open standards that better the industry? What's stopping us from open sourcing our code? Are we giving our users value, transparency, and control? Open up as much as you can as often as you can, and if anyone questions whether this is a good approach, explain to them why it's not just a good approach, but the best approach. It is an approach that will transform business and commerce in this still young century, and when we are successful we will effectively re-write the MBA curriculum for the next several decades!

An open Internet transforms lives globally. It has the potential to deliver the world's information to the palm of every person and to give everyone the power of freedom of expression. These predictions were in an email I sent you earlier this year (later posted as a [big post](#)) that described my vision for the future of the internet. But now I'm talking about action, not vision. There are forces aligned against the open internet --- governments who control access, companies who fight in their own self-interests to preserve the status quo. They are powerful, and if they succeed we will find ourselves inhabiting an internet of fragmentation, stagnation, higher prices, and less competition.

Our skills and our culture give us the opportunity and responsibility to prevent this from happening. We believe in the power of technology to deliver information. We believe in the power of information to do good. We believe that open is the only way for this to have the broadest impact for the most people. We are technology optimists who trust that the chaos of open benefits everyone. We will fight to promote it every chance we get.

Open will win. It will win on the Internet and will then cascade across many walks of life: The future of government is transparency. The future of commerce is information symmetry. The future of culture is freedom. The future of science and medicine is collaboration. The future of entertainment is participation. Each of these futures depends on an open internet.

As Google product managers, you are building something that will outlast all of us, and none of us can imagine all the ways Google will grow and touch people's lives. In that way, we are like our colleague Vint Cerf, who didn't know exactly how many networks would want to be part of this "Internet" so he set the default to open. Vint certainly got it right. I believe we will too.

The comments you read here belong only to the person who posted them. We do, however, reserve the right to remove off-topic comments.

8 comments:

...blogspot.com/.../meaning-of-ope...

6/9

12/22/2009

Google Public Policy Blog: The me...

Charles Darwin said...

Google incorporated in a private corporation engaged in free enterprise.

Google Inc. is Not a person can disappear in 5-10 years.

Witness Netscape, Sun Micro, Peoplesoft and DEC.

Gongle inc. competes in the free market, marketplace like all other companies.

Liberty and Freedom in a Free world will win in the end.

DECEMBER 21, 2009 9:21 PM

Paramendra Bhagat said...

A company that has become big needs a new ethos. Google as a company will stick around for a long, long time, but it is not at all a given that Google will continue to be on the bleeding edge. Will Google Wave end up being the last bleeding edge product Google gave to the world?

Staying open is your best bet to sta on the cutting edge for as long as possible.

DECEMBER 22, 2009 12:15 AM

David said...

I believe this is the way forward to making IT more accessible to everyone for nothing. IT is no longer a luxury but a necessity for our well being. Thank you Google.. :)

DECEMBER 22, 2009 1:15 AM

Lode said...

Call us when you open your search ranking algorithm, or the design for your servers. Until then this cherry-picked openness sounds a little bit hypocritical.

To quote the Register:

"The company claims to perform all kinds of magic to ensure that its data centers are the most energy-efficient computing houses around, saving Google millions of dollars per year on electricity. Yet the super-secretive ad broker continues to reveal precious little about these data center tweaks.

If Google is actually successful at lowering power consumption, then the secrecy makes sense. It sees the tweaks as a competitive edge over companies such as Microsoft and Yahoo!, which also consume immense volumes of gear and power.

At what point will Google decide that all service providers should have a chance to benefit from its knowledge for the betterment of mankind?"

DECEMBER 22, 2009 3:42 AM

vbhu said...

Great post Jonathan.

I liked the philosophy of openness. I have always wondered where does one draw the line

...[blogspot.com/.../meaning-of-ope...](#)

12/22/2009

Google Public Policy Blog: The me...

between the openness in a corporate environment - and your post here made me realize that it is a fine balancing act - and it keeps changing with the conditions in the market.

Ultimately, there are things which can be opened and some which can't. Most companies just hoard everything. Its refreshing to see Google open up so many things.

DECEMBER 22, 2009 6:25 AM

Karthik said...

This post has been removed by the author.

DECEMBER 22, 2009 8:58 AM

Karthik said...

Holy Google.

Overall, I am inspired. But there are SEVERAL holes in this position, as thy actions and words do not align.

Would thou "open" up thy search algorithm?

Should I assume, sometimes even GOD is imperfect.

Should I look at this differently? I seek enlightenment

DECEMBER 22, 2009 8:59 AM

zetrik said...

I tend to agree about Google's servers and their search algorithms being closed, it's not very nice.

I think that's how they survive for now though, if they released those informations, Microsoft would trample them.

They'll probably release some of that information once they have more other competitive products or when it's getting less risky to release it in general.

In the end, I think those principle Google is trying to define are challenged by reality. The current reality is that law is not in favor of those principles so Google has to do compromises.

DECEMBER 22, 2009 10:49 AM

Post a Comment

Links to this post

Open Google - Luca De Biase

Netizen: An Open Google

GiCi

Googland: [G] The meaning of open

Open Google

Create a Link

...blogspot.com/.../meaning-of-ope...

8/9

12/22/2009

Google Public Policy Blog: The me...

[Home](#)

[Older Post](#)

Copyright © 2009 Google Inc. All rights reserved.
[Privacy Policy](#) - [Terms of Service](#)