

TALKING PERFORMANCE WITH THE WORLD'S LARGEST ENCYCLOPEDIA

An Interview with Wikipedia's Sue Gardner

*In the eight short years since it went live, Wikipedia has become the Internet's standard encyclopedia and dominant general reference resource. It's a top five site in the world; some 320 million monthly visitors peruse more than 14 million Wikipedia articles. The English-language version receives millions of page views every hour, and peaked in March of 2009 at 11.913 billion page views for the month. **Benchmark** recently caught up with Wikimedia Foundation Executive Director Sue Gardner, who was selected as one of the **Huffington Post's** top ten "game changers" for 2009, people "who are harnessing the power of new media to reshape their fields and change the world." We talked with Sue about managing the world's largest encyclopedia, evolving the user experience (especially for editors), and measuring performance worldwide.*

Benchmark: We were surprised to find, and more than likely most people would be surprised to find the number of properties under the Wikimedia umbrella — a dozen all told. Wikipedia is certainly the best known, and must be the busiest of all your sites.

Sue Gardner: Yes, by far, by orders of magnitude. Wikipedia is the project that everybody knows and it's a project that's most used of all of them by far.

Benchmark: Can you quantify the traffic you get on Wikipedia?

Sue Gardner: Typically, how we like to measure ourselves is by unique visitors, because we think that's the closest approximation you can get of people served. Our purpose is public service — creating information and disseminating to big numbers of people. According to comScore, we typically receive about 320 million unique visitors a month.

Benchmark: What kind of infrastructure do you have to serve all those people?

Sue Gardner: Our primary data center is in Tampa, Florida. That's historically because the Wikimedia Foundation started in St. Petersburg. We also have a data center in Amsterdam. But that's it. We intend to build out a fully redundant data center elsewhere in the United States sometime over the next year; that'll probably be in the Washington, D.C. area. And we're also looking to build sets of caching servers elsewhere in the world just to increase

performance. Obviously working with Keynote is a helpful tool for that, because we would like more information about our global performance and we need to find out where we're weak before we can locate caching servers to strengthen up.

Benchmark: Okay, so you're looking for Keynote to help you characterize your performance worldwide?

Sue Gardner: That's our hope, yes. The upshot of our situation is that we have never really had consistent performance metrics that are broad and reliable. And so a relationship with Keynote is going to be really important for us for benchmarking how we're doing.

Benchmark: OK, so you have, on average, about 320 million unique visitors a month. Do you ever experience any big surges in traffic beyond that, in response to big news stories, for example?

Sue Gardner: Yes, huge surges in traffic on individual articles. Probably the first one for us was right after 9/11. Wikipedia started in January of 2001 and, of course, 9/11 was in 2001. Most of the major news and media sites went down; some of them for a couple of hours, some of them, if I recall correctly, for four or five days. I was working at the Web site of the Canadian Broadcasting Corporation at the time and we definitely went down, for a couple of days. At that time, Wikipedia didn't go down. We actually were fine, primarily because we were so small at the time that the traffic spike was relatively small compared to what it would have been today.

The New York Times wrote a story about the number of people at that time who visited Wikipedia for information. And we've seen that as a consistent trend. Every time there's a big news story, particularly a big international news story, we get massive spikes in traffic. When Hurricane Katrina hit, we got massive traffic. After the London bombing, we got massive traffic. The most recent example is when Michael Jackson died.

Benchmark: Right — everybody got hit when that happened.

Sue Gardner: We had a huge, massive spike. It's interesting, because Wikipedia started off as a conventional encyclopedia, with a lot of background information on topics that are of some sort of enduring permanent interest. But what's interesting is that

one of the key ways it's different from a traditional encyclopedia is, obviously, that it's continuously updated.

And so over time, as our base of encyclopedia material got fleshed out, the key changes and the key updates tend to be turning it into, in some ways, more of a news site — because it is super responsive and updated multiple times every second around the world.

Benchmark: That's an interesting evolution. So Wikipedia is now part of the real-time Web?

Sue Gardner: Oh, absolutely. I remember when Luciano Pavarotti died, I was watching CNN and I saw the news flash, breaking news — they were announcing it for the first time on CNN. So I immediately went to Wikipedia to update the article, and it had already been updated. I think it was updated before the news came across on CNN.

Benchmark: Pretty impressive, to be scooping CNN.

Sue Gardner: Yes, we don't break news because we don't allow original research, everything needs to be cited and sourced. But frequently we have it before people hear about it from other venues because there are so many Wikipedia editors — about 100,000 core editors around the world — that they'll hear something somewhere and they will update the article in question, usually before people get the information from another source. So increasingly, people are turning to us for that kind of thing.

Benchmark: It's interesting because Wikipedia is not something one would think of as a typical first choice for news.

Sue Gardner: Yes. But it's partly just Google too, right? If you search for Luciano Pavarotti, Wikipedia is one of the first results. If you search for Michael Jackson, we're one of the first results.

Benchmark: Very good. Okay, let's go back to the performance issue. What sorts of performance challenges do you face at Wikipedia?

Sue Gardner: I would say our biggest performance challenge is that we don't have consistent measurement of how we're doing. It's really that fundamental. So for example, our mission is to serve people around the world, and we tend to perform really well from a reach perspective in Europe and in North America. We tend to perform less well in Asia — there are probably a whole bunch of reasons for that. We don't know if the site is less responsive and slower if you're in Japan or China. So our major challenge really is getting in place those broad benchmarks of performance and seeing where we are weak. We don't have answers to those really

basic questions.

Benchmark: So then you don't have any kind of sense of whether you're losing people because of performance?

Sue Gardner: No, we don't. And clearly we need to. What's interesting about Wikipedia is that, as an organic mass-created endeavor, it's done fantastically well. It's in the top five Web site by unique visitors globally — it's number five. But it's gotten really far just through the organic effort of lots of different people. I believe it's now time for us to be more systematic about how we measure things and start explicitly, deliberately redressing the gaps, such as our lack of performance in Asia. That's what we're poised to do right now.

Benchmark: So the work that's pending with Keynote (at this writing) is really your first foray into actually measuring the performance of the site?

Sue Gardner: Yes, that's right. And it's really, really important to us. When Wikipedia first started it didn't matter. It was an experiment and if it was slow, if it was unresponsive, if it was down for a while, it wasn't a big deal. There have been a couple of times in our history where we've been down for fairly long periods of time, and that didn't matter when we were just a little Web site that no one really counted on. But over time, as it's gotten much, much more popular and therefore more influential, people have really come to count on it, and we have a much higher need to be accountable to people and to be there for them because they're relying on us.

One of critical challenges is the credibility of Wikipedia. We want people to trust it. We consider ourselves an educational endeavor and we want to partner with schools and universities and subject matter experts. In order to persuade people to trust us and to work with us in this important endeavor of making information freely available to everybody — in order for that to happen, we need to be perceived as reliable and responsive and accountable and credible.

And a piece of that is performance, because we know that an organization or a Web site's credibility will suffer, people won't believe it in as much, if it's unresponsive, if it's slow, if it seems broken, if it seems sluggish. So it's important for us to be seen as fast and responsive and reliable.

Benchmark: So the business ramifications for your performance deal primarily with credibility and with your ability to forge these partnerships?

Sue Gardner: Yes, that's a big piece of it. The credibility is one big

piece of it. But the other piece really is just an end in and of itself. We want to be there for people and that's our job. So it would not be okay if Wikipedia were down for several days or if it was really unresponsive for people. That would mean that we weren't serving our audience — we weren't doing our job.

Benchmark: Right. But the difference between delivering a page in two seconds and four seconds is not as critical for you as it is for the retailer that's trying to get somebody to buy something that they can buy in other places.

Sue Gardner: Well no, I think it's equally important for us. It's just important for different reasons. We are an international site, we have an international audience with international users. And we want to be equally responsive to all those people. I do think that a piece of our challenge in terms of reaching people in the parts of the world where we're not doing so well likely is linked simply to responsiveness and simple performance measures. One of the things that's good about Wikipedia is that it's a fairly fast site, because it's primarily text — it's not a heavy site.

An we pride ourselves on that. We want the output of Wikipedia to be available to people quickly. We want it available to them in a flexible way where they can view it on multiple devices and edit for multiple devices and use the content in a variety of different ways — and it's designed to make that possible. But basic things like responsiveness are really important because we know that they are drivers of credibility and usage.

Benchmark: Okay — social media. Wikipedia was really one of the first user-generated sites. Has the rise of other forms of social media impacted you in any way, or is there any cross pollination? Do you see any role for it?

Sue Gardner: Yes, definitely. But our core fundamental purpose is very different from most social media sites. We're very different from, for example, Facebook or LinkedIn. We're not a platform primarily intended for people to connect with each other.

Online community is not our goal, it's not our purpose. Our purpose is to collaboratively create a bedrock of free information for people, and the online community piece is how we get there. So in the Wikimedia properties, it's always been secondary. They aren't there to have fun. They aren't there to engage with each other, they aren't there to network with each other, they're there to build an encyclopedia.

Having said that, it is obviously the case that you need to have a healthy community of contributors and you need to provide all kinds of tools and functionality for them to interact successfully toward that end goal. We're really interested to see some of the

functionality that other sites have come up with on their own — we do think there's lots that we can learn from them.

For example, we're going to be launching a rate-this-page feature. We know that asking people if they think an article is high quality doesn't necessarily tell you if it's high quality, but it gives you some useful information. So we're going to be adopting things where we think they make sense for our audience and for our mission.

Benchmark: But the other types of tools and functionality you're talking about — they're more behind-the-curtain, for your content community. It's not visible or doesn't really have implications for the average person just coming to the site off the street, so to speak.

Sue Gardner: The Wikipedia editing community has developed all kinds of tools and functionalities to support its work which, as you say, aren't necessarily visible to the average reader — or if it is visible, it might be unclear what its purpose is or how it's used.

But one of our big strategic initiatives right now is trying to break down the wall between a reader and a contributor. We want to make it easier for people to contribute to Wikipedia. So we've launched a big usability initiative designed to make the MediaWiki software easier for folks to use in editing the encyclopedia. MediaWiki was developed many years ago when the entire Internet was difficult to use. It was normal for things to be difficult. Obviously since then, huge advances have been made in creating functionality that makes it easy for people to interact online. Yet MediaWiki is more or less the same in most ways as it always was.

So we did some really interesting user testing where people were asked to try to edit an article on Wikipedia and most of them recoiled in horror.

Benchmark: Yes.

Sue Gardner: Because — yes, have you tried?

Benchmark: Yes.

Sue Gardner: It's fairly horrifying, right?

Benchmark: Right.

Sue Gardner: So — we have a lot of usability challenges which we're attempting to overcome with various projects designed to simplify the interface and better support people. In general, we would like to see more people edit Wikipedia. The average Wikipedia editor is a young man, typically in Europe or in North

America, typically a grad student — 87% of our editors are male and they're on average 25 years old. Those are great people who form the core of our community. They're engaged in the world of ideas. They're thinking. They're reading. They're writing all the time.

Benchmark: But it seems like a rather narrow demographic.

Sue Gardner: The broader your community of contributors, the more diverse they are and the more different backgrounds and experiences and skills that they have, the better and the richer and the smarter and the more complete your final product is going to be. So we aspire to bring in a lot more women to edit.

And there's a huge bulge of the baby boomers who are starting now to retire. Those people have spent their entire lives learning things and understanding and we think that they can be persuaded to give back now — to contribute some of what they've learned to Wikipedia would be a fantastic thing for the Web site. And for the world.

We have a lot of work to do to make it easier for those people to contribute. So, yes, we do want to adopt stuff from social media where it makes sense for us.

Benchmark: So in terms of what the future looks like, it's looking a lot more user-friendly and a lot more inviting of participation from people.

Sue Gardner: We want to make it possible for people to edit Wikipedia who don't have 40 hours to figure out Wiki syntax. That need to invest in understanding the technology — that barrier — shouldn't exist. We have some challenges also around policies and so forth, because Wikipedia has developed over time, like I was saying earlier, like a newsroom. It has developed all kinds of policies around neutral point of view and citations and so forth and all that policy stuff is good bedrock material. If you're going to be an encyclopedia, there's a box of what an encyclopedia is and what an encyclopedia is not.

So policies are necessary, but at the same time, people shouldn't have to invest 100 hours in learning about Wikipedia before they can make a contribution to it because that's not the spirit of the endeavor. The spirit of the endeavor is that we all know something, we all know something about something and we should be able to contribute whatever it is that we know and we should be able to contribute in different ways.

The goal for us is to figure out how to engage more people and get more people involved. That's our primary interest.

Benchmark: So how far along are you on that project?

Sue Gardner: We've made quite a few strides. Our usability project started in early 2009. I believe it's on its third release and its work continues.

The multimedia upload project has just begun and we've also just begun a project which we're calling the Bookshelf Project. The purpose of it is to provide a set of invitational and training and coaching materials for people who want to contribute but don't know how.

We also do a lot of other initiatives like Wikipedia academies, which are face-to-face coaching activities. We try to get together a critical mass, a couple of hundred people and engage them and teach them how to use the encyclopedia and how to edit it. So we've done a lot of things but there's still lots more to do.

Benchmark: The more you describe it, the bigger it sounds.

Sue Gardner: Yes. It's about changing the world.

Benchmark: Right. Now, could we — can we anticipate seeing some changes perhaps to the way your homepage is set up, the types of things we see there to foster and encourage more involvement?

Sue Gardner: Okay. If you are a registered user, you can see next to your username, there will be little notation that says 'try beta.' And you can try the new skin, which the usability team has been working on, which is a simplified interface and a slightly prettier interface. That's our first crack at that.


We're also just at the very beginning of starting to think about flexible interfaces that are responsive to the culture and the country and the look and feel that is appropriate for each context. One thing I think might be a contributing factor to our lower performance in Asia is that Wikipedia looks so different in Asia from other Asian Web sites.

Currently, Wikipedia looks the same no matter where you are in the world — fairly gray and dry relative to other sites. I've spent a little bit of time on a couple of Korean sites and a couple of Japanese sites and they look very different from what the sites look like in Europe and North America. So that raises the question for me that, if I'm a Korean person and I go to Wikipedia, do I experience it as really different and not for me — from somewhere else? And I think that it's probably the case that I do. Most international organizations try to let the skin or the look and feel be dictated by the culture that's reading and using the site. So we're just beginning to explore what that might look like for Wikipedia.

Benchmark: Pretty much everything about Wikipedia is going under some kind of review, update, or enhancement.

Sue Gardner: We've been putting in place for the Wikimedia Foundation over the last year, the last 18 months, performance measurement of various kinds in terms of participation and reach and quality of content and so forth, and we're only really now getting to the technical performance and how we can ensure responsiveness globally and identify geographic bottlenecks, and that sort of thing. So Keynote, we hope, is going to be a big piece of that work for us.

Benchmark: Surely it will! We'll be looking forward to an update. Thank you for your time and insight.



Sue Gardner
Executive Director,
Wikimedia Foundation

Sue Gardner is executive director of the Wikimedia Foundation. Prior to joining Wikimedia in 2007, she ran CBC.ca, the award-winning Web site of the Canadian Broadcast Corporation. Earlier, she spent 10 years as a radio and television journalist, focusing on pop culture, social issues and media analysis, including a stint on the now-famous syndicated radio program, "As It Happens." Sue is a member of the Online News Association, the Society for News Design, Women in Film and Television, the Canadian Association of Journalists and Canadian Women in Communications.