CHAPTER 4

The Four Big Questions

As you saw in Chapter 3, there's a dauntingly long list of metrics you can track about your website. All of those metrics answer four fundamental questions about your site's visitors:

- What did they do?
- How did they do it?
- Why did they do it?
- Could they do it?

Without the answers to these four questions, you can't tell where to focus your efforts. You don't know what's working well or what's hopelessly broken, and you can't tell how to improve your business. Most companies that can't answer these questions try new strategies randomly, hoping to hit on one that works.

Armed with the answers to the four big questions, you can make informed decisions about where to focus your efforts. Would better design or better infrastructure improve your conversion rates? Are your competitors winning against you before prospective customers ever have a chance to see your offers? Do visitors prefer to search or browse? Are buyers leaving because they didn't like your prices, or because they came to the wrong site altogether? Are your biggest subscriber's users' complaints your fault—or theirs?

These are the kinds of decisions that make or break a business. To make the right decisions, you need to answer the four big questions.

What Did They Do?

The first question concerns what your visitors did, and it's answered through web analytics.

Visitors' actions speak for themselves. Analytics shows you what worked best, but won't tell you *why* something worked. The only way analytics can help improve your site is by showing you which content, messages, designs, and campaigns have the best

impact on your business. This is because analytics lacks context—it won't show you what was on a visitor's mind, or whether the site was fast during her visit, or how easy she found it to use.

Analytics was once relatively simple because web transactions were simple. Three things have changed in recent years that complicate matters, however:

Visitor interactions aren't just requests for pages
: Gone are the days of straightforward page-by-page interaction; instead, visitors stay on a single page, but interact with page components through DHTML, JavaScript, or plug-ins.

Visitor impressions start long before users visit your site
: To get a complete picture of a visit, you need to know what people are saying about you elsewhere that led a visitor to your site and set the tone for his visit.

Visits don't follow a set path
: Instead of browsing a predictable sequence of pages to arrive at a result, visitors explore in a haphazard fashion, often relying on searches or recommendations to move through a site.

While the basic building block of analytics is an individual visit, analysts seldom look at web activity with this much granularity. Instead, they look at aggregate analysis—metrics grouped by geography, demographics, campaigns, or other segments. You'll only look at individual visits when you're trying to reproduce a problem or resolve a dispute. Web analytics is more focused on large-scale patterns of interaction.

How Did They Do It?

The second question is all about usability. Most modern websites let users accomplish their goals in a number of ways. Visitors might search for a book or browse by a list of authors. They might click on text hyperlinks or the images next to them. They might take a circuitous route through the application because they didn't see the big button just out of sight on their screens. They might abandon a form halfway through a page. Or they might enter a zip code in the price field, then wonder why they can't complete their transaction.

Usability is a mixture of science and art. Web designers want their work to look fresh and innovative, but need to balance their desire for cutting-edge design with a focus on familiar affordances and button conventions that visitors understand.

Perhaps because of designers' eagerness to introduce fresh sites, they've given users many more interfaces to learn, making the Web less consistent even as its citizens become more savvy.

A profusion of browsers
: Opera, Chrome, Internet Explorer, Firefox, Flock, Camino, Safari, and others all render pages slightly differently. Sometimes the differences can be subtle; other

times, entire ads may not show up, JavaScript may not execute, or functionality may be severely impaired. You may not even be aware of these limitations.

A variety of devices
: Notebooks, desktops, mobile phones, and PDAs offer different controls. Some have touchscreens and new gestures, while others have numeric keypads. They all display different resolutions. We use them in many different surroundings—in a car, in bright light, or in a noisy room.

New interaction metaphors
: Drag-and-drop, wheelmouse-zoom, modal displays, and so on, are now possible with Flash, Java, and JavaScript.

Usability is most important when it affects goal attainment. It's the domain of Web Interaction Analytics (WIA).

The term "WIA" was coined by ClickTale, and is also variously referred to as in-page or behavioral analytics.

Why Did They Do It?

The third big question looks at consumer motivation. It's a far softer science than analytics. While analytics might offer tantalizing clues to why visitors tried to do something on your site, you're still left guessing. The only way to be really certain is to ask visitors directly.

On the Web, answering the question of "why" is done through VOC services—surveys that solicit visitor opinions, either within the site itself or through a third-party survey. Respondents are recruited (by mail, for example) or intercepted as they visit the site. They are then asked a series of questions.

With the emergence of online communities, website analysts have new ways of understanding consumers. Web marketing is much more of a two-way conversation than print or broadcast media. As a result, analysts can talk directly with their market and understand its motivations, or analyze its sentiments by mining what communities are discussing.

Could They Do It?

There's no point in asking what users did if they couldn't do it in the first place. Answering the fourth question means measuring the health, performance, and availability of the application. As with web analytics, this task has become more complicated in recent years for several reasons.

An application is no longer made from a single component
: Your application may rely on plug-ins and browser capabilities. It may pull some parts of a page from a Content Delivery Network (CDN). It's dependent on stylesheets and JavaScript functioning correctly. It may be delivered to a mobile device. And it may incorporate third-party content such as maps. Even within a single web server there may be various operating systems, services, and virtual machines that make it hard to properly instrument the application.

You're present on sites you don't control
: You may have a Facebook group or Twitter profile that drives users to your main site. Your site may incorporate third-party SaaS elements: a helpdesk and FAQ service, a portal for job seekers, and an investor information page with stock prices. Even if your own site works well, your users' experience depends just as much on these external content sources.

Some of your visitors are machines
: If you have any kind of API—even an RSS feed—then you need to worry about whether scripts and other websites can correctly access your application. Even if you're just being indexed by search engines, you need to help those engines navigate your site.

Measuring site health from the end user's point of view is the discipline of EUEM, which is the logical complement to web analytics. It's often measured in terms of performance and availability, from both individual and aggregate perspectives.

There are two major approaches to capturing user experience: *synthetic testing*, which simulates user visits to a site, and *real user monitoring* (RUM), which watches actual users on the site itself. Together, they provide an accurate picture of whether visitors could do what they tried to do on the site.

RUM data may be collected through many of the same mechanisms as WIA, and some analytics or RUM products offer WIA functionality, such as replay and click overlays.

Putting It All Together

Answer these four big questions, and you're well on your way to improving your website, because you have a comprehensive understanding of how all its pieces fit together. For example:

- If user experience is suffering, you can increase capacity, reduce page size and complexity, deliver content via CDN, replace components that are too slow, and so on.
- If you're failing to meet agreed-upon service level targets, you can mitigate SLA arguments and reduce subscriber churn by contacting your users to tell them you know they're having problems, and that you're going to improve things.

- If your conversion rates aren't high enough, you can try to better understand what users want or where they get stuck.
- If visitors take too long to complete transactions, or don't scroll down to key information, you can make the site more usable so they can easily find what they're looking for.
- If you're converting the visitors that make it to your site, but aren't getting enough traffic, you can focus on communities and word-of-mouth marketing rather than site design.

Ultimately, by answering the four big questions you close the loop between designing and deploying a website, and adapting quickly. You make mistakes faster and find the right way to connect with your market.

Analyzing Data Properly

Before we look at how to answer these four questions in more detail, we need to talk about data analysis. Many of the monitoring technologies we'll cover rely on statistics to analyze millions of pieces of data about thousands of visitors quickly and easily. If you don't look at that data with an analytical eye, you can easily be misled.

Web analytics, EUEM, VOC, and WIA provide a tremendous amount of raw information about your website. You need to analyze and communicate it properly if it's going to have an impact on your organization. That means comparing metrics to other things, segmenting measurements into useful groups, and using the right math for the job.

Always Compare

Data analysis is all about comparisons. You should always talk about data with words that end in "er": better, faster, stickier, heavier, weaker, later. You can compare groups of visitors, periods of time, or versions of content. You can also compare yourself against your competitors.

To compare, you need a sense of what "normal" is—a baseline or control group. The "always compare" rule means your first job will be to establish baselines for key metrics like conversion rate, visitor traffic, web page performance, visitor satisfaction, email inquiries, and call center volume. Only then can you make useful comparisons in future reports.

Segment Everything

Whether you're trying to judge the effectiveness of a marketing campaign, the cause of a problem, the usability of a page, or the importance of customer feedback, you need to segment your measurements into several groups.

Grouping data into manageable, meaningful clumps is essential. As humans, we naturally try to cluster data into segments. The challenge is in knowing which of the thousands of possible segments is most likely to yield the right data.

Sometimes, the data will segment itself in obvious ways—by geography, browser type, referring site, carrier, and so on. Other times, you'll have to create your own segments along which to analyze the measurements you collect. No matter what you're trying to do, having several segments to compare shows you which is best and which is worst. From there, you can start to fix things.

Segmentation applies everywhere. If you're trying to resolve a performance problem, your first question will be whether there is a particular segment for which the problem is more common: is it offshore visitors? Are all the affected visitors on the same broadband carrier? Is it always the same page, or the same server? Similarly, if users aren't seeing part of the page, is the problem related to age groups? Browser types? Screen resolutions?

Don't Settle for Averages

British Prime Minister Benjamin Disraeli once said, "There are three kinds of lies: lies, damned lies, and statistics" (or perhaps he didn't, though Mark Twain attributed this famous quote to him). The *really* compulsive liars, however, are averages. Averages are misleading, often concealing important information.

Consider, for example, a class of 20 children who are 5 years old. The average age in the room is five. When a 90-year-old grandparent comes to visit, the average age climbs to 9. This is misleading: we'd prefer to know that the average age is five, with one outlier, or that the most common age in the room is five.

For meaningful measurements, insist on histograms (frequency distributions) and percentiles. A histogram is simply a chart of how many times something happened. Figure 4-1 shows the age histogram for the classroom we just visited.

Histograms are particularly useful for analyzing quantitative information (like page load time), as they show us at a glance how many measurements fall outside a certain range or *percentile* of the samples we collected.

When we talk about "the 95th percentile of latency," we mean the delay that 95 percent of visitors experienced. Figure 4-2 shows the distribution of performance across a day's page requests, alongside the average latency and the 95th percentile latency.

Percentiles are good because they let us quantify the impact of a problem. If the 95th percentile page latency suddenly jumps to 16 seconds, as it has in this example, we now know that we have a serious problem affecting a broad swath of visitors (five percent of them, in fact). What do those page requests have in common? Are visitors all asking for the same page? Are they all from the same city? Are requests only from new visitors?

This is where diagnosis begins—by identifying an issue and segmenting to find out what's common across all those who experienced the issue.

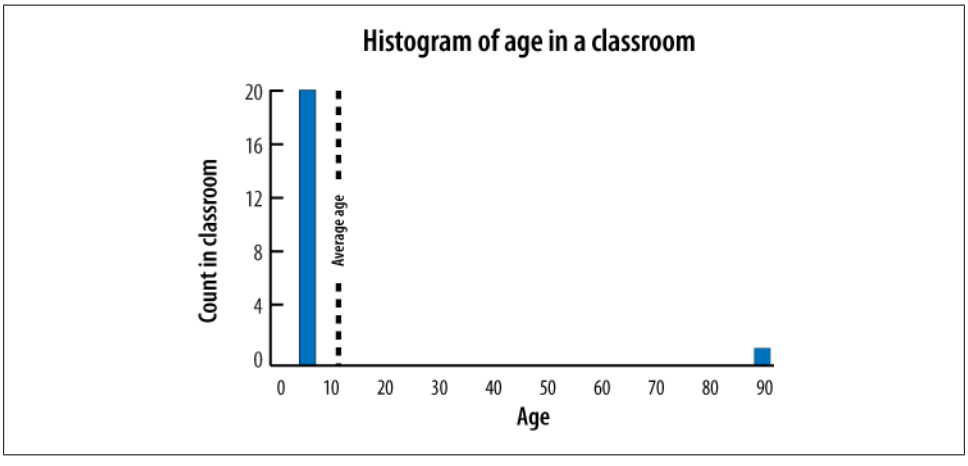

Figure 4-1. A histogram displaying age in a classroom paints a more accurate picture than an average would

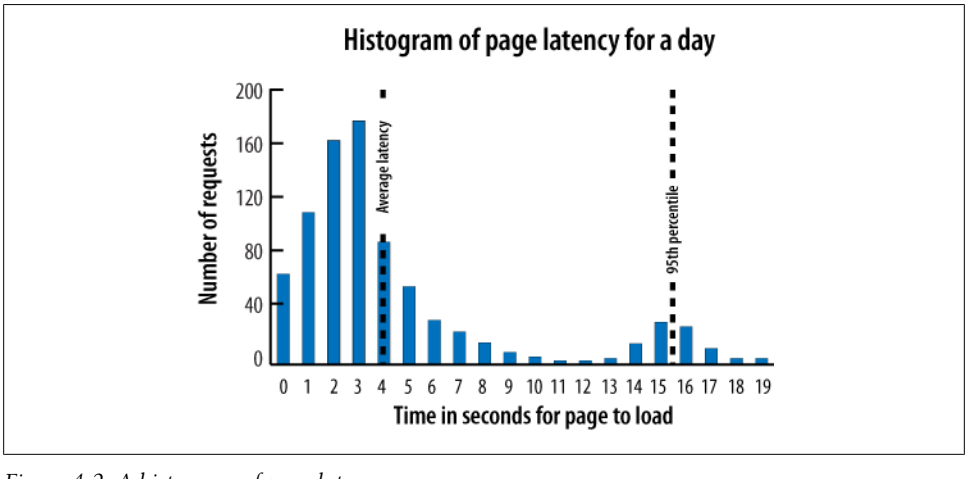

Figure 4-2. A histogram of page latency

A Complete Web Monitoring Maturity Model

As organizations improve their web visibility, they transition through several stages of maturity. You won't achieve the same maturity across all four of the big questions at once. You may already have mature web analytics, for example, but only primitive insights into web performance.

Level 1: Technical Details

At the first level of maturity, monitoring is done for technical reasons. The organization hasn't yet realized that the web application is important to its business strategy, and results are used only to ensure that systems are functioning and to detect egregious errors. If you're at this level, you're probably asking questions like:

- How many hits does my site get each day? Is that number growing?
- What percentage of HTTP GETs to the site receive an HTTP 200 OK response?
- How slowly do my servers respond to the tests I'm running?
- Does my website work on different web browsers?

There's nothing fundamentally wrong with these questions, but on their own, they don't relate to the business because they don't relate to a particular business objective.

Consider, for example, a server that gets a thousand requests. That traffic could all be coming from one energetic user, or it could be a thousand one-page visitors who didn't find what they wanted. Neither is good. Or consider a server that suddenly responds slowly to a test. Unless you know how many users were on your site and what they were doing, you don't know whether the slowdown affected your business.

Level 2: Minding Your Own House

As an organization engages its customers, it first focuses on the things it can control. This leads to questions about the *effectiveness* of the website, such as:

- Does the website convert visitors?
- Is it fast and available when people need it?
- Why do visitors come to the site?
- What are the most common ways people use the site, and which of those lead to good outcomes (goals and conversions)?

These are the key metrics that organizations need to use to make sure the site can do what it needs to: convert visitors. They're a great start.

Because they don't look beyond the borders of the website to the Internet as a whole, however, organizations at this level are still inward-facing and somewhat myopic. They also fail to consider long-term metrics like the lifetime value of a visitor, loyalty, and capacity trends.

Level 3: Engaging the Internet

With your own house in order, you can now turn your attention to the rest of the Web. The main focus here is driving traffic to your site through marketing campaigns, paid

search, and organic search. The site engages visitors in other ways, too, forcing you to ask new questions:

- How are people finding the site, and what's working best?
- Are the other components of the Internet, such as online communities, APIs, web services, and CDNs, functioning properly?
- What's the Internet saying about the organization, and how is that affecting traffic rates?
- Are different segments of the Internet showing different levels of usability? Are foreign-language browsers having to scroll down more? Are older visitors with larger font sizes seeing less of certain pages?
- Are there trolling and spamming problems?

To deal with the flood of measurements generated by additional visitors, the organization starts to analyze data differently. For example, it uses software to automatically baseline what "normal" is, pinpoint exceptions, and predict seasonal or growth trends.

Level 4: Building Relationships

The next stage of maturity turns visits into long-term relationships. In the earlier stages of the maturity model, each visit was treated as a unique interaction. As your organization matures, these interactions are stitched together into a user record, and analytics starts to converge with CRM.

The central focus is now a visitor's lifetime relationship with your site. You learn how often he visits the site, what triggers those visits, and what outcomes happen.

Interactions happen beyond the website. Customer loyalty cards show store visits, calls to the call center are tracked by phone number, and so on. The relationship becomes a two-way engagement in which you reach out to visitors who have opted in to mailing lists or who subscribe to RSS feeds.

A relationship-focused organization cares about longevity, loyalty, and capturing the most complete view of its visitors:

- How often do visitors return, and what's their lifetime value to the organization?
- How steep is the learning curve, and how long does it take a new visitor to use 80 percent of the site's features?
- How do word of mouth and the community affect loyalty?
- Which modes of interaction (web, online community, mail, phone, in-person) are most common, and what causes visitors to switch from one to another?
- How many visits does it take before a reader comments?
- How willing are my customers to have me communicate with them?

- What's my reach (ability to get to customers) and how likely are they to act on what I send them?
- What performance or uptime expectations—or implied SLA—do my I have with my power users?

Stitching together individual visits, both offline and online, is a major challenge for any organization. If you're at this stage, you need to handle personal information carefully so you don't run afoul of privacy laws or put your organization at risk.

Level 5: Web Business Strategy

The final stage of maturity occurs when the business makes the Web an integral part of its product and marketing strategies. In the earlier levels of the maturity model, the Web is an "online branch office"—it often has its own profit and loss, and is run like a separate business.

But the Web touches every part of our lives. It's on our phones and in our cars. It's quickly becoming the dominant medium for music, video, and personal messaging. It's how employees work remotely. It's many people's primary source of day-to-day knowledge. The Web's effects reach far beyond a storefront, affecting everything from how companies develop new products to how they support existing ones.

When companies finally embrace the Web as the disruptive force it is, they stop becoming reactive and start becoming opportunistic. Instead of reacting to what web monitoring tells them, they start to wonder whether the Web makes new things possible. When you reach this stage of maturity, you'll be asking questions like:

- Can I develop and roll out products more quickly, and with greater accuracy?
- Can I reduce support costs and make customers self-supporting?
- Can I tie my services to real-world components (for example, linking my CRM portal to a GPS so salespeople can optimize their travel)?
- Can I move into other media or gain presence elsewhere?
- How does the Web remove barriers to entry in my industry? How does it let me erect barriers to new competitors?
- Can I engage target markets in product design and research?

One clear sign that a company views the Web as strategic is that it starts to include Web metrics in employee performance reviews and corporate goal-setting. The Web isn't a side business anymore: it *is* the business.

The Watching Websites Maturity Model

Table 4-1 shows the various levels of maturity. We'll return to the table as we look at how to answer the four big questions.

Table 4-1. Stages of watching websites

Component		Maturity Level 1	Level 2	Level 3	Level 4	Level 5
Focus		*Technology: Make sure things are alive.*	*Local site: Make sure people on my site do what I want them to.*	*Visitor acquisition: Make sure the Internet sends people to my site.*	*Systematic engagement: Make sure my relationship with my visitors and the Internet continues to grow.*	*Web strategy: Make sure my business is aligned with the Internet age.*
Who?		Operations	Merchandising manager	Campaign manager/SEO	Product manager	CEO/GM
Web analytics		Technical details: Page view, hits. Focus on operation of the infrastructure, capacity, usage.	Conversion: How many of the visitors complete the goals I intend?	Traffic: How does the Internet learn about me, encourage visits to me?	Relationship: How often do buyers return? What's the lifetime value? Where else do customers interact with me?	Strategy: How can I combine my brick-and-mortar and web businesses? How does the web change my company?
	Synthetic	Availability and performance: Checking to see if the site is available from multiple locations, and reporting on performance.	Transactions and components: Multi-step monitoring of key processes, tests to isolate tiers of infrastructure.	Testing the Internet: Monitoring of third-party components and communities on which the application depends.	Correlation & competition: Using the relationship between load and performance; comparing yourself to the industry and public benchmarks.	Organizational planning: Using performance as the basis for procurement; uptime objectives at the executive level; quantifying outages or slow-downs financially.
EUEM	**RUM**	Errors and performance: Counting hard errors (404, 500, TCP reset) and end user performance grouped by customer or visitor segment.	Analytics integration: Tying user experience to business outcomes within the site to maximize conversions; identifying "soft errors" in transactions.	All components, all functions, automation: Watching content from third-party sites and user actions within a page; automatically forming baselines and diagnosing exceptions.	SLA, CRM: Using RUM information as the basis for SLAs, problem resolution, and release management.	Integrated user engagement: Measuring user experience across call centers, in-person encounters, web transactions, etc., as a single entity; impact of user experience is quantifiable.

Component	Maturity Level 1	Level 2	Level 3	Level 4	Level 5
VOC	"Contact us" buttons and on-site feedback; emphasis on satisfaction.	Surveys within the site via opt-in invitations; emphasis on loyalty.	Engaging the public Internet (chatrooms, social sites, etc.) and analyzing key topics and discussions; emphasis on word-of-mouth and virality.	Customer collaboration in product and service design; user engagement; emphasis on lifetime value creation, giving the user a sense of ownership.	Consumer feedback tied in to corporate planning through quantitative analysis of VOC and community data; customer as a collaborator in the growth of the company.
WIA	Click diagrams showing "hot" areas on key pages.	Segmentation of user actions (scroll, drag, click) by outcome (purchase, abandonment, enrollment).	Segmentation by traffic source (organic search, campaign) and A/B comparison; visitor replay.	Learning curve analysis; comparison of first-time versus experienced users; automated A/B testing of usability.	Product specialization according to usability and user groups; usability as a component of employee performance.