

EVALUATING THE IMPACT OF YOUR WEBSITE

A guide for United Nations and CGIAR
Centers to evaluate the usage,
usability and usefulness of their
websites



VERSION 1.0 – NOVEMBER 2007



ict-km

Collaborate, Create, Communicate

A PROGRAM OF THE CGIAR 

EVALUATING THE IMPACT OF YOUR WEBSITE

A guide for CGIAR Centers to evaluate
the usage, usability and usefulness
of their websites

M.BUONAIUTO

*Head of Digital Unit
United Nations CGIAR
Rome, Italy*

ict-km

Collaborate, Create, Communicate

CGIAR

The Consultative Group on International Agricultural Research (CGIAR), established in 1971, is a strategic partnership of countries, international and regional organizations and private foundations supporting the work of 15 international Centers. In collaboration with national agricultural research systems, civil society and the private sector, the CGIAR fosters sustainable agricultural growth through high-quality science aimed at benefiting the poor through stronger food security, better human nutrition and health, higher incomes and improved management of natural resources. www.cgiar.org

ICT-KM

The ICT-KM Program of the CGIAR promotes and supports the use of information and communications technology (ICT) and knowledge management (KM) to improve the effectiveness of the CGIAR System's work on behalf of the poor in developing countries. For more information see www.ictkm.cgiar.org

Correct citation: Buonaiuto 2007. Evaluating the Impact of Your Website: a guide for UN CGIAR centers to evaluate the usage, usability and usefulness of their websites. ICT-KM Program of the CGIAR, Rome, Italy.

This work is licensed under the Creative Commons Attribution-Noncommercial-Share Alike 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-sa/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.



You are free:

- to Share — to copy, distribute and transmit the work
- to Remix — to adapt the work

Under the following conditions:

- Attribution. You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).
- Noncommercial. You may not use this work for commercial purposes.
- Share Alike. If you alter, transform, or build upon this work, you may distribute the resulting work only under the same or similar license to this one.
- For any reuse or distribution, you must make clear to others the license terms of this work. The best way to do this is with a link to this web page.
- Any of the above conditions can be waived if you get permission from the copyright holder.
- Nothing in this license impairs or restricts the author's moral rights

Rights owner: ICT-KM Program of the CGIAR (ictkm@cgiar.org)

Editor: SCRIPTORIA Writing and Training (www.scriptoria.co.uk)

Graphic designer: Paolo Pontieri

Table of contents

Acknowledgements	5
Foreword	7
Overview	9
The Website as Part of a Center's Strategy	10
Elements to Consider when Planning a Website or Website Section	11
<i>Step 1: Set the Focus and Scope</i>	<i>11</i>
a) Define the purpose	11
b) Define and prioritize the target audience	13
c) Setting objectives	14
<i>Step 2: Select the Channels and Packaging</i>	<i>14</i>
<i>Step 3: Decide on Access</i>	<i>15</i>
<i>Step 4: Promote the Website</i>	<i>15</i>
<i>Step 5: Monitor and Evaluate Results</i>	<i>16</i>
1. Measuring Usage	17
The Theory	17
<i>Two Web Analysis Technologies</i>	<i>17</i>
Web log analysis	18
Web-based analytics	20
<i>Choosing the Right Analytics for the Job</i>	<i>20</i>
Putting it into Practice	23
<i>An Example of Web Log Analysis</i>	<i>23</i>
Logging	23
Logging methods	23
Logging and performance	23
Logging and local data protection law	24
<i>Understanding Web Analytics</i>	<i>24</i>

Filtering log files	25
<i>Analysis of Trends</i>	26
General overview of results	27
Specific analyses of Web campaign results	27
<i>Template for a Web Analytics Report</i>	29
Analysis of users	29
Analysis of content	32
Analysis of thematic areas of a website	35
Referrers	37
Statistics on RSS (Really Simple Syndication)	42
Further Reading	43
2. Measuring Usability	45
The Theory	45
<i>When to Measure Usability</i>	46
Putting it into Practice	51
<i>A Hypothetical User Test</i>	51
<i>Polls and Surveys</i>	57
3. Measuring Usefulness	59
The Theory	59
<i>Types of User Surveys</i>	60
Putting it into Practice	62
<i>A Hypothetical Usefulness Survey</i>	62
Identifying the survey subjects	62
Drafting the questionnaire	64
Administering the survey	68
Analyzing and interpreting the results	69
Appendix A: Problems and Limitations of Web Log Data	70
Appendix B: How to Use Google Analytics	75
Appendix C: Website Performance Test	80
Appendix D: Glossary of Terms	82





Foreword

The plethora of information available on website evaluation is almost overwhelming, a situation that is compounded by the availability of an extensive array of evaluation tools capable of generating different types of statistics, tables and graphs. This makes it difficult to decide which website elements should be evaluated and how best to use the information generated to improve the impact of your website.

CGIAR Centers, Programs and Projects all produce websites. But what is the return?

This document is the first step towards the creation of an easy-to-use guide for evaluating the impact of websites. Customized especially for CGIAR websites, the guide's format is designed around the 3 Us approach: assessing the Usage, Usability and Usefulness of a website and its content. For each area it provides a background to the evaluation process in the 'The Theory' section, and guidelines on how to apply that knowledge in 'Putting it into Practice'; the latter also provides practical examples. Although how you apply the guidelines will vary according to your particular situation, and resources and time you have available, this guide provides a sound starting point for evaluating your website.

This guide is not a technical guide, but aims to bring people with communications and information management expertise together with those who have the technical expertise in a common platform and using common language for planning and evaluating websites.

This guide is presented as Version 1, for two reasons. Firstly, the frequent changes in website technology options will need to be incorporated at a future point. Secondly, the guide should be viewed as an evolving document, with more practical examples added over time. In particular, it would be useful if users of this guide could share with us their evaluation results, their interpretation of those results, and the decisions made as a

consequence. These can then be incorporated into future versions of the guidelines.

After we have learned from the evaluation of other websites, we will develop a new guide to help plan websites for impact.

We have created a space on CGXchange (www.cgxchange.org) called “Websites & Impact” where all CGIAR staff can contribute ideas and experiences that will enrich this document and help develop the next version.

The ICT-KM Program welcomes feedback from the information, communication and technical communities or anybody who can use or contribute to this guide.

The information in this guide is one of several outputs of the ICT-KM Program’s Web Content and Usage Analysis Project, which was tasked with developing strategies and tools for marketing Web content and evaluating usage. Special thanks go to Joanna Kane-Potaka of Bioversity International for leading these efforts.

*Enrica Porcari
CGIAR Chief Information Officer
Leader, ICT-KM Program*

Overview

This guide provides an introduction to measuring the impact of your website based on the 3Us—usage, usability and usefulness.

Each section presents a general discussion of how to measure each characteristic—‘The Theory’—and, examples of how to apply this knowledge—‘Putting it into Practice’—and suggests approaches and resources for further learning.

It must be stressed that using Web analytics, usability testing and feedback surveys, to name the principal tools of these respective undertakings, is definitely a non-trivial pursuit. It is not something that can be immediately mastered simply by reading a short guide. Scores, if not hundreds, of books have been written on these subjects. Nevertheless, a journey of a thousand miles begins with a single step. Hence this brief introduction.

It also must be emphasized that none of these activities will be useful unless they occur in the context of what the owners of the website want the website to achieve, both in general terms and in terms of any specific aspects. As Jason Burby of the Web business consultancy ZAAZ, Inc. said in a recent ClickZ article, the three reasons that Web analytics most often fail companies are:

- The goals for the analytics program are ill- or un-defined.
- Information generated from analytics projects is not effectively shared within the organization.
- Companies don't take action based on the data they collect.

The key then is for the website to be purpose-driven. Only then will evaluating the impact be meaningful. This means maintaining a constant creative process involving both those who measure the impact of a website and those who possess the tools, access and authority to improve it.

The Website as Part of a Center's Strategy

The Web is the primary information source and outlet for millions of people in today's connected world. A center's website can serve many different functions, including contributing towards the center's strategies in the following areas:

Information and knowledge sharing. The interactive and collaborative options of Web 2.0 tools provide new opportunities for two-way sharing of information and knowledge.

Marketing research projects' outputs. The website is a key tool in promoting the uptake of research results.

Public awareness. The website and Web tools can complement and strengthen other parts of special campaigns.

Branding. A website is both highly visual and rich in content and plays a central role in efforts to consistently reflect the desired image of a center and its key messages.

Public relations. The website is a key reference source used by people to learn about a center.

Resource mobilization. The website can be a tool for resource mobilization, for example by giving donors access to information, projecting an image and messages important for fund raising, and even as a mechanism for receiving donations.

Marketing. The marketing strategy of a center can draw upon all of the strategies mentioned above to market the center as a whole or specific outputs.

Web publishing and Web development should always be done with a clear purpose.

The development of a website, or special sections of the website, needs to be driven by what the website or website sections are expected to achieve—website development is not just a technical project.

Elements to Consider when Planning a Website or Website Section

Key steps to consider when developing a website or website section are listed below. This process will need to involve technical staff, management groups, editors and communication experts if it is to deliver better products and more useful evaluation reports.

Step 1: Set the Focus and Scope

a) Define the purpose

A website can play many different roles in a center's strategy, as listed above, ranging from public awareness to knowledge management and resource mobilization. As a result, defining the purpose of a website can be quite complex. In some cases the website as a whole may have one purpose while a particular section of the website may have another very different and specific purpose.

So, it is recommended that the 'purpose' be defined for:

- the website as a whole
- each section of the website
- any change in the architecture or visual design of the website.

Defining the purpose of the website as a whole As an example, Bioversity might define the purpose of its website to include the following:

To be a good source of scientific information and knowledge on agricultural biodiversity, allowing visitors to the website to access and share information.

This would be important for:

- substantiating the center's scientific credibility with target audiences (as defined in the center's strategy), especially the scientific community, industry players internationally and donors.
- assisting the center's scientists in their work.
- assisting other scientists (especially in developing countries) in their work.

To be a basic source of information on agricultural biodiversity for the general public.

This would be important for:

- substantiating the center's reputation internationally with target audiences that are not scientific researchers, for example the media, policy-makers, community groups and some donor groups. (Note that many of the individuals in these groups may have a scientific background although they may be working in a non-research capacity.)
- increasing awareness and understanding of the key issues among the general public and the center's non-scientific audiences (as mentioned above).

Providing corporate information about the center.

This would be important for:

- substantiating the center's professional corporate credibility with all of its target audiences.
- providing specific corporate information needed by different target audiences (for example, giving donors access to annual reports).

- strengthening corporate branding, including portraying the image (personality traits and values) of the center to target audiences.

Defining the purpose of a section of the website The purpose of a section of a website could, for example, be to stimulate the interaction between agriculturalists and health workers in a campaign on biodiversity and nutrition. This might be done by developing e-discussion and interview sections on the website.

Defining the purpose of a change in the architecture or visual design of the website The purpose of making changes to the architecture or design of the website could be, for example:

- to give more prominence to certain sections (by changing the navigation structure).
- to strengthen the image of the center as an organization devoted to science or poverty alleviation (by changing the look and feel of the website).
- to assist in more targeted searching for scientific information (by introducing new features for searching the website).

b) Define and prioritize the target audience

When defining the purpose of a website, it's usual to identify the target audience. The target audience may be defined quite broadly. For example, the target audience for a CGIAR center may include:

- Investors
- Research collaborators (ranging from NARS to local community groups and global/regional organizations)
- The scientific community
- Users of the research results (e.g. farmers, NARS, policy-makers, international organizations)
- Influencers of the above groups (e.g. media and general public and key individuals/organizations)
- Developed and developing country groups (for all of the above).

It is useful to prioritize these groups and identify specific target audiences in line with a center's strategies. For example:

- For the Public Awareness strategy, the priority target audience for a campaign on climate change may be policy-makers in key developed countries;
- For the Resource Mobilization strategy, the priority audience may be six donors who are being specifically targeted for funds.

c) Setting objectives

Objectives for the website will be much more specific than the purpose. Objectives should have SMART characteristics. They should be:

- Specific—about what is to be achieved
- Measurable—so that progress towards the objectives can be quantified
- Achievable—the objectives should be achievable within the agreed timeframe
- Realistic—the objectives should be achievable given the resources and time available
- Time bound—a specific deadline should be set for achieving the objectives

Objectives may be quantitative or qualitative and will guide the collection of statistics and information for evaluation.

Statistics about previous operations, and information and benchmarks relating to other similar websites, may also help set realistic and achievable objectives.

Step 2: Select the Channels and Packaging

There are a variety of features and sections on the website that can be used to channel information and different ways each of these features or

sections can be packaged (designed). Some examples of website channels and packaging include:

- themed Web page
- newsletter
- e-discussion forum
- reference section
- success stories and case studies
- media release
- GIS or similar visual mapping presentation

Step 3: Decide on Access

Access may be chosen consciously or by default. Some questions to consider are:

- Is access to information or a particular section restricted or freely available?
- Is there easy access for areas of poor connectivity?
- Are there any hindrances to access through, for example, sign-in processes, use of technologies that are not accessible to all computers/systems?

Step 4: Promote the Website

Planning promotional activities should be integral to the development of new websites, Web sections, information or features of a website. Some questions to consider here are:

- Which content or features will be promoted?
- Will links be requested from other websites and other products?
- Will efforts be made to contribute information and links to self-editing sites such as Wikipedia, Wiktionary and blogs?
- What metadata needs to be created for the content to be crawled

by search engines, social websites and tagging systems (e.g. keywords, XML sitemaps)?

- Which traditional promotional means will be used, for example bookmarks, articles in newsletters, displays?

Step 5: Monitor and Evaluate Results

How the website will be monitored and evaluated to check whether or not the expected outcomes were achieved and why.

If any aspect of the planning process is overlooked the website may lack the focus necessary to achieve the center's objectives. This is especially true when objectives have not been clearly defined; the danger then is that the website will lack clarity of purpose and direction, and will gradually become an amorphous collection of unrelated pieces of information. Refocusing such a site is an enormous task. It is far better to have a clear plan from the beginning and to constantly monitor progress towards the objectives.

The following sections of this guide will help you monitor and evaluate the performance of your website—but only if you have defined what that performance should be in the first place.

1. Measuring Usage

The Theory

Measuring usage is important to a website manager for several reasons. First, basic measures of the number of users, visits and resources accessed can provide a rough estimate of whether the site is perceived as useful by users. Clearly, if nobody uses the site, it isn't doing much good.

Second, usage data can be used to improve the site's usability and usefulness. Analyzing key processes, such as finding or downloading static or dynamic content, buying or downloading something, or contacting the site operators with a question, and breaking down these processes into the steps needed to complete them, can show whether the process was completed, or, if not, at what step it was abandoned. In operational terms the process is defined by the pages or applications that need to be accessed along the way. Experimental changes can then be made at the point where most users abandoned the process. Usage data will show whether or not the completion rate increased after the changes were made. This process of continuous improvement can be repeated until real improvement occurs.

Finally, usage data can measure the effectiveness of attempts to promote and market the website. For example, changes in usage of the website, or of particular pages, can be correlated with efforts to increase their usage. For instance, the number of people logging in to a Web page could be measured for periods before and after an email was sent out promoting content on that page.

Two Web Analysis Technologies

Analysis of website usage is most often carried out with one of two technologies, either Web log analysis or Web-based analytics. Each technology has advantages and disadvantages.

Web log analysis

Web log analysis is the oldest major method of measuring Web traffic. Web logs were originally set up on Web servers to measure their technical effectiveness. Later, analysis of Web log data provided important usage information. Web logs are limited, however, for several reasons, including the type of data that can be collected as part of the Web server operation and the limitations imposed by standards. Web log analysis is still the most widespread method of measuring Web traffic, however, and because of this there are many techniques and a large body of literature for using and interpreting Web log data.

The data collected by Web servers that are relevant to traffic analysis include: resource requested, date, client IP (Internet Protocol) address, referrer, URI (Uniform Resource Identifier) query, HTTP (Hypertext Transfer Protocol) status, time taken, and cookie. Cookies, identifiers on the Web client, are particularly useful. If the presence of a cookie can be matched with registration information collected from a user, it is then possible to cross-tabulate user characteristics with activity on the website. If each user uses only one computer, which, in turn, has a unique IP address, it is possible to infer a great deal simply from associating client IP addresses with resources requested. However, proxy servers, the use of DHCP (Dynamic Host Configuration Protocol) by Internet Service Providers and several other issues make this difficult. This is discussed in more detail in Appendix A.

Nevertheless, Web log analysts have managed to spin a large number of key performance indicators out of massaging these data, such as overall traffic volumes, number of visits, ratio of new to returning visitors, percentage of new visitors, page stickiness, referring domains, search keywords and phrases, percentage of visits of a certain duration, average pages per visitor, average visits per visitor, entry pages and content viewed, the proportion of heavy users, average time spent on site, key conversion rates, and activity of 'retained' visitors.

Conversion funnels, or process measurement tools, analyze how a particular visitor accesses resources, beginning with the page where they enter a website and ending with an action such as purchasing a product. The metrics required to perform such an analysis include not only Web log data but also data from databases that underlie the website (for example, records of user logins) and data provided by visitors, such as query string values.¹

There are downsides to using Web logs, however. Web logs can become large and hard to manage. Running custom analyses can take time. Teasing out just the data you want requires significant analytical skill. If you are interested in data from many sites, the Web log analysis needs to combine data from the logs of each server. This is often impractical. Furthermore, getting consistent results from analyses of the same log data from different Web analysis software packages has often been problematic, even when the only difference has been shifting to an upgraded version of the same package.

Web logs also do not record when a user re-uses a page cached on the client computer, which can lead to misleading measures of usage. Perhaps most importantly, Web logs do not record data about individual users. Using cookies can improve this, but then a database has to be set up to store and provide a means for retrieving data about individual users. This isn't all that hard, but somebody has to do it. Web log analysis programs also have difficulty interpreting dynamic pages, such as those generated by Content Management Systems; the content viewed by the user depends on the request they make and the permissions they have in the system.

¹ Eric T. Peterson, *Web Analytics Demystified: A Marketer's Guide to Understanding How Your Web Site Affects Your Business*. Portland: Celilo Group Media and CafePress, 2004. pp. 20–21.

Web-based analytics

In recent years, Web-based analytics using JavaScript tags on pages have become more popular. This technology does not store usage data in server's Web logs but in a special computer run by an application service provider (ASP). The data is sent from the client to the ASP, as well as being sent to the Web server. Thus, the ASP can combine data from sessions on many servers.

A large publisher, for example, can combine data about all its publications, or a global organization like the CGIAR can combine data from all of its centers.

The main drawbacks of Web-based analytics are that vendors each have their own set of metrics, which may not be exactly what you want. The standard set of metrics is mostly defined by the needs of large online retailers rather than by the needs of organizations such as the CGIAR. While the basic statistics and process analyses may be easy to convert to useful information, real customization is more difficult than when you have local control of the data and the analysis package. Also, there are issues with ownership of the data and what might happen if you switch vendors or encounter different analytical formats. Also, depending on the size of your operation, Web-based analytics may be expensive.

Google has recently started offering Web-based analytics, named, strangely enough, Google Analytics. The service is free for websites with less than 500,000 page views per month. Many CGIAR websites fall in this category. CGNET uses this service for its own site, in addition to the AWStats Web log analysis available to all websites hosted at CGNET. (See Appendix B for information on Google Analytics.)

Choosing the Right Analytics for the Job

The most important thing about measuring anything is knowing what you want to find out. This is probably even more important for a typical CGIAR center website, because the goals of such a site may be quite different

from the goals of websites for which most books about Web analytics are written, or for which most Web analytics software is designed.

Web analytics are obviously crucial to retail sites that are trying to sell products over the Web. They are also important to Web-based publishers and nonprofits who aim to raise money from Web users. Although it is possible to posit such activities for CGIAR center websites, these are not the usual purposes.

Why is it important for a CGIAR center to increase traffic to its website over time? There are two reasons. The first is because website traffic is an indicator of general interest in the center and this should at least hold steady. Second, the website may be a means to promote the center, a campaign, or a point of view. In this case, increasing traffic, particularly to any specific pages being promoted, is important.

Often, website owners attempt to 'drive' traffic to the site by sending out press releases, exchanging links with other websites, advertising, or direct marketing. Measuring the results of these efforts is important.

At a minimum, there are probably four basic reasons for paying attention to Web analytics at a CGIAR center.

The first is to make sure that the website is operating properly. For example, Web analytics record the number of errors a website generates. So if a website designer has redone a Web page and left a link pointing nowhere, clicks on that link will be recorded as errors, and the Web statistics will identify the link to be fixed. Other measures, such as which operating system or browser clients are using, can be helpful in redesigning a site when deciding, for example, what size Web page to offer, and whether special efforts have to be taken to accommodate different browsers.

The second reason for using Web analytics is to examine whether users are coming to the website for the reasons that the website's owners think

they are. The analysis of queries from search engines that send users to a website can be quite revealing. For example, the second most popular query for IFPRI (after “ifpri”) is “green revolution.” Looking at the rankings of the website pages can also be helpful in finding out why users are coming to a website.

The third reason for using Web analytics is to determine whether a website is easy to navigate. It is fairly safe to assume that any website operator wants a website that makes it easy for users to find what they want. This is, of course, the province of usability testing, which we discuss in Section 2, but Web-based analytics can also support usability testing. This kind of analysis allows the Web designer to see whether usage is flowing in the desired direction. A low percentage of clicks on an advertisement may indicate that there is room for improvement. In this case, the current advertisement could be replaced with a new one and the clicks could be measured again. By a gradual process of trial and error, the percentage of clicks may be increased. In effect, this repeated process of measurement constitutes a quantitatively based usability study.

The fourth reason is to assess whether the information on the website is useful. We will discuss usefulness in Section 3, but it is generally agreed that some kinds of usage constitute broad indicators of usefulness. One such indicator is repeated visits to the website by the same user, the assumption being that if the website is useful, it will be remembered and used again. A second indicator is the amount of time users spend on the website. Clearly, if a user leaves the website in less than 30 seconds, a statistic measured by AWStats for example, it is hard to believe that the user found the website useful.

Conversely, websites with longer average periods of use are often praised for being ‘sticky,’ which is to say they have content that users take the time to absorb. Of course, external factors can affect this statistic. Slow website performance, for example, will inflate the average period of use. Poor navigation may also increase the period of use. Some of these explanations can be ruled out by examining performance statistics as well

as by examining the time spent on specific pages. Another indicator of usefulness is whether the entry and exit pages for a given session are the same. The assumption here is that using a website usually involves going into it more deeply than one page.

Putting it into practice

An Example of Web Log Analysis

As part of the CGIAR system-wide contract, CGNET performs regular website analyses of CGIAR websites hosted at CGNET. These are available at <http://webusage.cgnet.com> for the 13 centers and Secretariat websites hosted by CGNET. Examining these pages will provide examples of regularly collected statistics of interest.

The monthly histories provide a good indication of how traffic is increasing or decreasing over time. Results for the current month and previous year are displayed, month by month.

Logging

Logging methods

The standard method of logging is to write data in a file on the server hard disk. But other methods are possible and often useful. The log data can be redirected into a database. The advantage of this is that log information from different servers can be sent to the same database. This makes it easier to compare different servers and services. Plus, if every session is associated with a unique session ID, it is possible to reconstruct single user sessions.

Logging and performance

It is useful to log a lot of information but this hampers performance as all the data must be stored and written to the hard disk. Quite often the logging process is a performance bottleneck. In the first month

after a launch of a Web application it may be necessary to log a lot of information but, step-by-step, this should be reduced.

Logging and local data protection law

In some countries it is important to check the kind of information that is logged and who will have access to it. For example, in Germany it is not permissible to store data that identifies individuals unless they have given special permission.

Understanding Web Analytics

'Web analytics' is the analysis and reporting of website traffic.²

Saner and Kanjorski (Gartner Group) state that:

"Web analytics uses a variety of data and sources to evaluate Web site performance and visitor experience, potentially including usage levels and patterns at both an individual and an aggregate level. Data and sources may include click-stream data from the Web server log, Web transaction data, submitted data from input fields on the Web site and data in the Internet customer repository. The goals are to improve site performance, both from a technical and content perspective, enhance visitor experience (and thus loyalty), contribute to overall understanding of customers and channels, and identify opportunities and risks."

An important definition of Web analytics is given in Web Analytics Demystified ³:

"Web analytics is the assessment of a variety of data, including Web traffic, Web-based transactions, Web server performance, usability studies, user submitted information and related sources to help create a generalized understanding of the visitor experience online."

² <http://www.answers.com/web+analytic?gwp=12&method=2>

³ Eric T. Peterson, Web Analytics Demystified: A Marketer's Guide to Understanding How Your Web Site Affects Your Business. Portland: Celilo Group Media and CafePress, 2004.

The underlying assumption is that the information gained will be used to improve the overall quality of the visitor experience.

The reader will immediately recognize that this definition is very broad and includes a multitude of data sources not mentioned here.

Analysis of statistics is no substitute for initial analysis of audience requirements. However, any Internet marketing strategy should include analyzing the results of Web analytics and taking any necessary action, bearing in mind that data from Web logs must be used cautiously. Server log data records the operations of machines and transactions (individual requests for files), not of people, and thus has significant limitations (see Appendix A for an explanation). Any analysis should take a holistic approach considering the log data together with what is known about user behavior.

Responses to online surveys and, in general, any other data provided voluntarily should also be used with caution because informants are 'self-selected' and thus possibly not representative of broader audience(s).⁴ Nevertheless, data generated in these ways can provide useful information about the typology of users interested on specific Web assets, such as training materials or publications.

Filtering log files

Before analyzing log files, it is necessary to create a sound statistical sample. This can be done by filtering the log files.

Log files register all requests made by the Web server. So, it is important to remove certain items ('statistical noise') from the initial set of requests.

This can be done by following two simple steps:

⁴ <http://depts.washington.edu/pett/papers/arthritis/Ramey.doc>

Step 1: Exclude robots and crawlers from the analysis.

Search engines, worms and viruses merely scan websites seeking data (content, emails, etc.). They must be excluded from the analysis because they do not involve any human interaction.

Step 2: Identify and exclude internal staff from the analysis.

Internal staff should not be included in the main analysis but dealt with in a separate short report.

These two steps will reduce the size of the sample and the sample will be 1) coherent and 2) statistically valid.

The absolute numbers and trends in results from a set of filtered log files will be different from those in results from a set of log files that have not been filtered.

Recently, Google Analytics introduced a new service that filters data properly, excluding internal traffic and specific IP addresses.

Analysis of Trends

Evaluation of Web metrics should focus on trends and not on absolute numbers.⁵

Standard reports can be generated regularly every 3, 6, or 12 months, or as required by the management and staff. The reports should present the results of the analysis and highlight trends that are emerging, without focusing on absolute numbers. They could be organized as follows:

⁵ "Don't Obsess Over Absolute Numbers, Focus on Trends", WebTrends report. (http://www.ksg.harvard.edu/itforms/stats/Trending_final.pdf)

General overview of results

- **Analysis of visitors** This presents statistics about users who accessed the website in the reporting period. The visitors can be classified by geographic location, relationship to the organization, etc. This section could include predictions of trends for the next reporting period.
- **Analysis of content** This presents an analysis of the data on access to Web pages and thematic areas of the website, and makes recommendations to facilitate decision-making.
- Other sections of the report can present analyses of statistics related to downloads of Web assets, such as PDF and PowerPoint files, or access to databases and responses to online surveys, under headings such as 'Analysis of publications' or 'Analysis of training materials'.
- **Analysis of internet referrers**—the URL of the Web page from which the visitor came—can provide useful data on the route by which visitors arrived the website. Strategic marketing campaigns based on referrers can significantly change the typology and number of visitors to a website. An analysis of accesses from other CGIAR centers or FAO websites could also be included in this section.
- The report may also indicate the ranking of the website in search engines and the main search phrases visitors use to bring them to the website.
- A section on 'Analysis of entry pages and exit pages' can indicate how users' comprehension of the content offered by the website was affected by the layout of the pages they visited.

Specific analyses of Web campaign results

- Key performance indicators (for example, Return On Investments for training material in terms of cost per download of PDF files, number of visits to new training modules).
- Comparison of actual with expected results.

A template for such reports is presented on the following pages. Reports can be enriched with additional statistics, such as new vs. returning visitors, depth of visit, browser capabilities, but you will need to consider the reason for including each additional metric and the implications of any analysis.

Template for a Web Analytics Report

Analysis of users

Access statistics

Access statistics show how many times visitors accessed the website and include accesses, sessions, sessions by one-time users, sessions by repeat users, etc.

Table 1. Accesses to website.

Accesses	Month 1	Month 2	Month 3	% change	
				Month 1/Month 2	Month 2/Month 3
Total accesses	75,755	96,028	94,431	27%	-2%
Total sessions
Sessions by one-time users
Sessions by repeat users
Total session users
Total duration of all sessions
Average accesses per session
Average sessions per user
Average session duration

Interpretation: The last two columns ('% change') represent the **trends** and are the most important columns to consider; absolute numbers often are too imprecise to provide statements and recommendations for the readers.

User sessions

Table 2. User sessions.

Session users	Month 1	Month 2	Month 3
One-time users	79.93%	80.71%	83.45%
Repeat users	20.07%	19.29%	16.55%
Two-time users
Three-time users
Four-time users

Interpretation: This table indicates the loyalty of visitors. One-time users generally represent 70–80% of the users of almost all websites. The table can be used, for example, to check whether ‘pushing services’ to promote content, such as newsletters or RSS feeds, are increasing user loyalty.

Visitors

Table 3. Visitors.

Website section	Number of visitors			% change in number of visitors	
	Month 1	Month 2	Month 3	Month 1/Month 2	Month 2/Month 3
About Us	5,905	10,708	17,400	81.33	62.49
...					

Interpretation: This table should not be considered the same as the 'Page views' table (see below), even though, as may happen when the Web page content is successfully marketed, the number of visitors may increase in parallel with the 'page views' metric and vice versa.

Table 4. Geographic location of visitors.

	Geographic location	Page views	Visitors	Total time spent by all visitors (dd hh:mm:ss)	Average time spent per visitor (hh:mm:ss)
1	United States	68,775	10,825	2d 20:53:31	00:00.23
2
3
4	...				

Interpretation: It is likely that the largest numbers of website visitors will be from North America, Europe and other developed countries, simply because the numbers of Web users are much higher in those countries than in less developed countries. However, given that developing countries are a major client group for the CGIAR centers it is important to monitor their usage of our websites. If use of the website by developing countries is very low, you should investigate why this might be, for example your site may have many 'heavy' graphics or your website has not been actively promoted in those target countries, and seek solutions to increase use by target countries. Comparing use before and after changes will indicate whether your changes have encouraged greater use among the geographical group you wish to target. You should also routinely monitor trends in use by geographic location so as to pick up any changes early on and identify possible causes.

Comment: The information provided by analyzing Table 4 could be taken further by looking at 'page views per visitor by country' (see Table 5). The

analysis could be limited to countries with more than a certain number of visitors. In the example below, the limit is 80 visitors.

Table 5. Page views per visitor by country.

Geographic location	Page views	Visitors	Total time spent by all visitors (dd hh:mm:ss)	Average time spent per visitor (hh:mm:ss)	Page views per visitor
Ethiopia	1,820	100	00:44:00	00:00:26	18.2
South Africa	1,541	86	00:53:50	00:00:37	17.9
United States	68,775	10,825	2d 20:53:31	00:00:23	6.4

Interpretation: The analysis and discussion of data in this table should focus on the last column, 'Page views per visitor', which indicates the extent of interest your site has for different geographic groups of user. It is important to monitor changes over time, especially if you have added significant new content to your website or have actively promoted particular parts of the website. If the number of page views per visitor from developing countries increases, this should be highlighted because these countries are a primary target for our information.

Analysis of content

Page views

Table 6. Page views.

Website section	Page views			% change in page views	
	Month 1	Month 2	Month 3	Month 1/Month 2	Month 2/Month 3
Publications	22,188	33,131	42,571	49.31	28.24
...					

Interpretation: This table presents the number of page views of particular sections of the website. The number of page views before and after changes made to the website—updates or new content—during the reporting period can be compared. If the updates or new content were promoted by newsletters, RSS, or other means, then the number of page views would be expected to increase. If the number of page views does not rise, there are two issues to consider:

1. These statistics may be an aggregate for many pages within a section. As a result it can be difficult to see the effects of changes to only some of the pages within the section, for example new publications listed in the publications catalogue or pages promoted through a newsletter or RSS feed. In this case, further investigation is required, such as checking the number of visitors that accessed the new content.
2. The updated or new content has not been promoted sufficiently. This can be confirmed by looking at the table for the number of visitors. In this case, a Web campaign should be initiated to promote the new content and an ad hoc Key Performance Indicator created to measure page views in the next reporting period.

Time spent on the Web site

Table 7. Total time spent in different website sections.

Website section	Time spent in website section (dd hh:mm:ss)		
	Month 1	Month 2	Month 3
Publications	4d 07:08:58	8d 13:52:37	8d 23:47:51
...			

Comment: With the coming of Web2.0, Flash, etc., the ‘page views’ metric is no longer sufficient to determine actual usage of a website. It is preferable to measure the time users spend in a specific section of the website.

The ‘time spent’ metric is affected by the timeout settings of the log analyzer.

*Most accessed pages***Table 8. Most accessed Web pages, ordered by visitors.**

Web page or address	Page views	Visitors	Page views per visitor	Current ranking	Past ranking
Homepage	55,358	29,297	1.88	1	1
/Publications/pubfile.asp (details about publications)	2	3
Employment page	3	2
Search page	4	5
Article in online magazine	5	12
...			

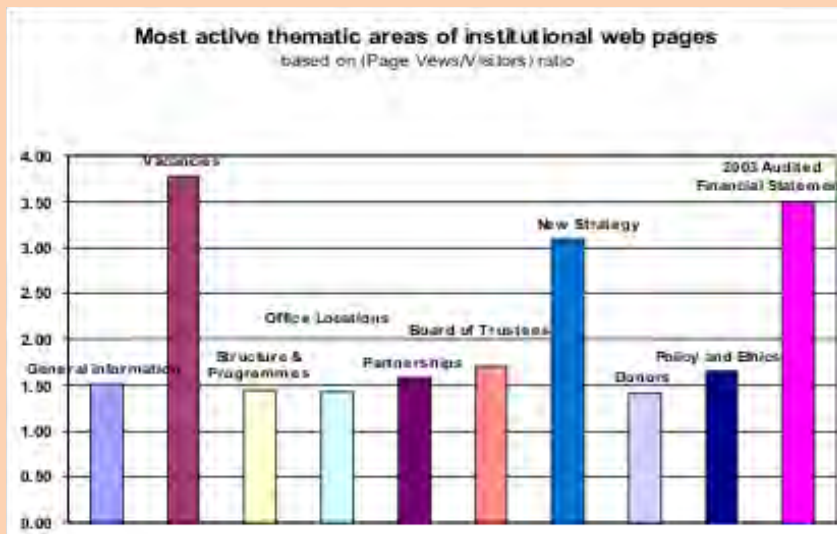
Comment: The employment vacancies page and the homepage are often the most accessed pages.

This table can be used to seek answers to specific questions such as: Did the resources invested in a specific Web page bring more users to the page? Is a specific database accessed by users as expected? To do this, you will need to look at changes in numbers of visitors and page views over time, and not just changes in ranking.

Analysis of thematic areas of a website

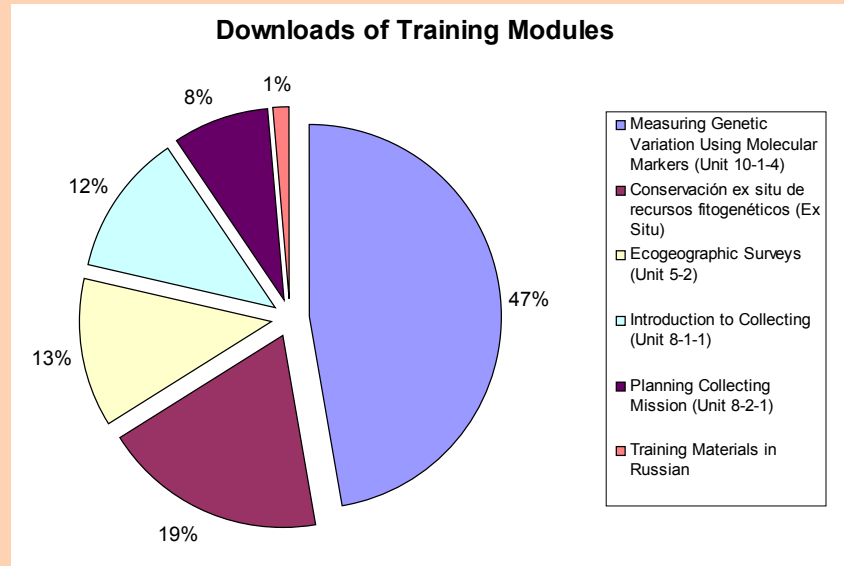
Analysis of the metrics presented in the previous sections may warrant a more critical look at thematic sections of the website (scientific areas, programmes and network sub-Web pages). Such an analysis could be done in collaboration with the directors of these areas and should be combined with online surveys, analysis of the use of online databases (looking at downloads of training materials, for example), etc. Graphs such as those shown in Figures 1 and 2 may help managers to plan changes to their website.

Figure 1. An example of an analysis of page views and visitors for thematic sections of a website.



Summaries of findings, such as the following characterization of users who download training materials, may also be useful to managers.

Figure 2. An example of an analysis of downloads of training materials



Example: Cross analysis characterization of users who download training materials

A typical user who downloads training materials has accessed the website from an educational or government institution and prefers to download materials in PDF⁶ format, although PowerPoint presentations are also acceptable. A typical user wants materials on "XXXX", but also finds other topics, such as genetics (genetic diversity, genetic resources), molecular biology, conservation and biotechnology, interesting.

The majority of users prefer English (63.50%) or Spanish (30.91%). A few users prefer French (3.38%), Chinese (1.14%), Arabic or Russian (both 0.53%).

⁶ Post hoc analysis shows that the PDF format is the exclusive choice for at least 50% of users.

The training material downloaded is used for research, training in the classroom and to acquire knowledge (self-learning and consultation).

Analysis of interest in vacancy announcements

Vacancy announcements represent one of the most visited areas of many websites.

Table 9. Downloads of vacancy announcements.

Title	Downloads
Crop Professional Manager	567
....

Referrers

A referrer is a website that a visitor looked at immediately before getting to the current site and is represented by a URL. Referrers include, for example, search engines, university websites, scientific websites and commercial domains.

Search engines

Table 10. Number of visitors referred by search engines.

Search engine	Number of visitors		% change
	CURRENT PERIOD	PAST PERIOD	
Google	14,042	12,321	13.96
Yahoo	1,737	1,121	54.95
...	

Interpretation: The number of visitors arriving at the website from the different search engines is an indicator of how 'visible' the website is in the search engine results. Large changes, especially reductions in the number of referrals from a particular search engine, should be investigated urgently.

Search engine optimization techniques can improve the visibility of a website in one or more search engines.

Table 11. Number of visitors referred by other websites.

Referrer	Number of visitors		% change
	CURRENT PERIOD	PAST PERIOD	
www.fao.org (FAO Web site)	986	430	129.30
www.economist.com	563	960	-41.35
bioversity-pa.grinfo.net	
www.futureharvest.org
www.agroecology.org
www.sciencecases.org
en.wikipedia.org
...

Interpretation: This table can show the benefits of linking with others' websites. If an external website has added a link to your website, this should result in increased referrals, which should show up here. This can help demonstrate the value of time and efforts invested in working with third-party websites. For example, if the center has created an article about one of its mandate crops in Wikipedia one would expect an increase in referrals from that source. If there is no significant increase, one should assess whether to continue investing time and efforts in that referrer. Another use of this table is to assess whether readers find the content interesting enough to create links and references in their website and/or blogs.

*Top 50 search phrases***Table 12. Top ranking search phrases.**

Search phrase	Number of searches
chenopods	1,556
Vigne	1,034
Plantas alimenticias	886
coconut tree	578
picea abies	524
vanuatu	432
plantas frutales	343
bancos de germoplasma	323
pinus pinea	321
world vision	312
cgjar	305
...	...

Interpretation: The list of most commonly used search terms indicates those topics that visitors to the website are most interested in. This could be used to direct efforts to enhancing those areas that are of most interest to the visitors, or to adjust the messages delivered by the site if many of the popular search terms are not directly relevant to the site's core content.

Entry and exit pages

Table 13. Top ranking entry pages.

Entry page	Sessions
Homepage	86,206
Publication details	...
Vacancies Web page
...	

Table 14. Top ranking exit pages.

Exit page	Sessions
Homepage	55,034
Publication details	...
...	

Interpretation: The metrics for entry and exit pages can lead to wrong assumptions if interpreted out of context. For example, the Homepage is generally the top ranking entry page, but it's also the top ranking exit page and the most visited page. So, a useful metric could be to compare the number of users who visited the homepage with the number who visited the homepage and left straight away.

Cross analysis could be performed by filtering the referrers appropriately, as in the following example:

Impact of FAO links to the website

Filtering by referrer, in this case from the domain www.FAO.org, would show up the pages of interest to visitors arriving from that referrer. It would also allow analysis of the country of origin of users arriving from that referrer, as well as other information about the users, including specific pages within the referrer's domain that provided the link to our website. Such analyses would highlight the impact of references and links to the website in the referrer's website, as well as indicating gaps where there were fewer referrals.

Path analysis

Path analysis is the process of tracking the sequence of pages followed by a visitor, particularly in relation to an event, such as requesting a newsletter or downloading a publication. The object is to understand how users find the information they are seeking and what features of the website help them to do so.⁷

Analyzing the path followed by users as they browse the website can be deeply influenced by, for example, Javascripts, framesets and AJAX, so the results are often not reliable.

Statistics on RSS (Really Simple Syndication)

Statistics on RSS feeds are tough to catch properly and are best collected using online services such as Feedburner (<http://www.feedburner.com>), Clicky (<http://getclicky.com>) and StandardStats. (See <http://www.pronetadvertising.com/articles/web-analytics-rss.html> for more details.)

⁷ See Wikipedia (http://en.wikipedia.org/wiki/Path_analysis_%28computing%29)

Further Reading

Using Web analytics effectively requires much more information than can be provided here. Some good introductions to the subject are:

Peterson, Eric T., "Web Analytics Demystified: A Marketer's Guide to Understanding". Portland: Celilo Group Media and CafePress, 2004.

Sterne, Jim, "Web Metrics: Proven Methods for Measuring Web Site Success". New York: John Wiley & Sons, 2002.

Tyler, Mary E. and Jerri Ledford, "Google Analytics". Indianapolis: Wiley Publishing, 2006.

2. Measuring Usability

The Theory

Jakob Nielsen, perhaps the world's most recognized expert on website usability, offers the following definition of usability, which I quote at length:

*Usability is a **quality attribute** that assesses how easy user interfaces are to use. The word "usability" also refers to methods for improving ease-of-use during the design process. Usability is defined by five quality components:*

- **Learnability:** How easy is it for users to accomplish basic tasks the first time they encounter the design?
- **Efficiency:** Once users have learned the design, how quickly can they perform tasks?
- **Memorability:** When users return to the design after a period of not using it, how easily can they reestablish proficiency?
- **Errors:** How many errors do users make, how severe are these errors, and how easily can they recover from the errors?
- **Satisfaction:** How pleasant is it to use the design?

*There are many other important quality attributes. A key one is **utility**, which refers to the design's functionality: Does it do what users need? Usability and utility are equally important: It matters little that something is easy if it's not what you want. It's also no good if the system can hypothetically do what you want, but you can't make it happen because the user interface is too difficult. To study a design's utility, you can use the same user research methods that improve usability.*

Why Usability is Important

*On the Web, usability is a necessary condition for survival. If a website is difficult to use, people **leave**. If the homepage fails to clearly state*

*what a company offers and what users can do on the site, people **leave**. If users get lost on a website, they **leave**. If a website's information is hard to read or doesn't answer users' key questions, they **leave**. Note a pattern here?*

There's no such thing as a user reading a website manual or otherwise spending much time trying to figure out an interface. There are plenty of other websites available; leaving is the first line of defense when users encounter a difficulty.

Source: Jakob Nielsen, "Usability 101: Introduction to Usability," Jakob Nielsen's Alertbox, August 25, 2003. (<http://www.useit.com/alertbox/20030825.html>)

When to Measure Usability

People leave sites and pages that they find difficult to use. This is the key to one method of measuring usability, which is to measure how many users abandon the website at each stage of an important process. The five components of quality discussed above can be the basis for measuring usability and for an ongoing program to continuously improve usability.

Ongoing measurement of usability is not usual, however.

The literature on usability emphasizes over and over again that user testing should be done during the website design process and the earlier the better. User testing is part of the iterative design approach that has come to be the most accepted process for website design. A recent extensive review of usability literature by the National Cancer Institute's Communication Technologies Branch in the U.S. Department of Health and Human Services rated the iterative design approach highest, both in terms of importance and strength of supporting evidence¹.

¹ Sanjay J. Koyani, Robert W. Bailey, Janice R. Nall, et al., Research-Based Web Design & Usability Guidelines. Washington, D.C. U.S. Department of Health and Human Services, 2004. Section 1:2.

The best approach to improving the usability of a website, therefore, is to begin user testing immediately, if this was not done when the site was originally designed. Alternatively, user testing can be part of website redesign. Nielsen recommends the following process for user testing:

How to Improve Usability

*There are many methods for studying usability, but the most basic and useful is **user testing**, which has three components:*

- *Get hold of some **representative users**, such as customers for an e-commerce site or employees for an intranet (in the latter case, they should work outside your department).*
- *Ask the users to perform **representative tasks** with the design.*
- ***Observe** what the users do, where they succeed, and where they have difficulties with the user interface. Shut up and let the users do the talking.*

It's important to test users individually and let them solve any problems on their own. If you help them or direct their attention to any particular part of the screen, you have contaminated the test results.

To identify a design's most important usability problems, testing five users is typically enough. Rather than run a big, expensive study, it's a better use of resources to run many small tests and revise the design between each one so you can fix the usability flaws as you identify them. Iterative design is the best way to increase the quality of user experience. The more versions and interface ideas you test with users, the better.

User testing is different from focus groups, which are a poor way of evaluating design usability. Focus groups have a place in market research, but to evaluate interaction designs you must closely observe individual users as they perform tasks with the user interface. Listening to what people say is misleading: you have to watch what they actually do.

When to Work on Usability

Usability plays a role in each stage of the design process. The resulting need for multiple studies is one reason I recommend making individual studies fast and cheap. Here are the main steps:

- 1. Before starting the new design, **test the old design** to identify the good parts that you should keep or emphasize, and the bad parts that give users trouble.*
- 2. Unless you're working on an intranet, **test your competitors' designs** to get cheap data on a range of alternative interfaces that have similar features to your own. (If you work on an intranet, read the intranet design annuals to learn from other designs.)*
- 3. Conduct a **field study** to see how users behave in their natural habitat.*
- 4. Make **paper prototypes** of one or more new design ideas and test them. The less time you invest in these design ideas the better, because you'll need to change them all based on the test results.*
- 5. Refine the design ideas that test best through **multiple iterations**, gradually moving from low-fidelity prototyping to high-fidelity representations that run on the computer. Test each iteration.*
- 6. Inspect the design relative to **established usability guidelines**, whether from your own earlier studies or published research.*
- 7. Once you decide on and implement the **final design**, test it again. Subtle usability problems always creep in during implementation.*

Don't defer user testing until you have a fully implemented design. If you do, it will be impossible to fix the vast majority of the critical usability problems that the test uncovers. Many of these problems are likely to be structural, and fixing them would require major rearchitecting.

The only way to a high-quality user experience is to start user testing early in the design process and to keep testing every step of the way.

Source: Jakob Nielsen, "Usability 101: Introduction to Usability," Jakob Nielsen's Alertbox, August 25, 2003. (<http://www.useit.com/alertbox/20030825.html>)

Having said this, it is easy to imagine the reader reacting with despair. It is very common for usability testing to be avoided because of a lack of resources and heavy deadline pressures, even though there is a basic agreement that it should be done. User testing doesn't have to be expensive, however, and it doesn't have to delay a site's implementation. As Nielsen writes:

Yes, big computer companies have been known to invest in million-dollar usability labs. Yes, experienced usability professionals are very highly paid. And yes, large-scale user testing to compare multiple design alternatives across several countries can cost \$200,000 or more. Ouch.

*But most **everyday usability projects are cheap**. Small companies don't need labs; you can run user tests in a spare conference room. Rather than hiring expensive usability professionals with ten years' experience, you can teach existing staff how to conduct studies. And, even though international studies are great, you don't start there: just spend a few days testing five domestic customers.*

Even with a budget of \$200, you can do usability. The methods are incredibly flexible and scale up or down according to circumstance. On average, best practices call for spending 10% of a design budget on usability. That's a cheap way to ensure that you spend the remaining 90% correctly, rather than blow your budget on an unworkable design.

Case study reports are typically issued by companies that have followed the entire user-centered design process to the letter, starting with field studies. These descriptions of the great and the good can be daunting for projects with smaller budgets and tighter schedules.

Usability need not be grandiose. The simplest user testing method I recommend takes three days, but even faster tests are possible—

especially if you use methods like paper prototyping, which lets you crank through new design iterations in a few hours.

*One of the main benefits of letting user research drive design is that you don't have to spend time on **features that users don't need**. Early studies will show you where to focus your resources so that you can ship on time.*

*Finally, usability can save time by helping you quickly **settle arguments** in the development team. Most projects waste countless staff hours as highly paid people sit in meetings and argue over what users might want or what they might do under various circumstances. Instead of debating, find out. It's faster, particularly because running a study requires only one team member's time.*

Source: Jakob Nielsen, "Misconceptions About Usability," Jakob Nielsen's Alertbox, September 8, 2003. (<http://www.useit.com/alertbox/20030908.html>)

Putting it into practice

A Hypothetical User Test

As discussed above, the particular circumstances the operators of a website face determine how and when to test users. The optimum time is during the initial design of the website or during a redesign. As an example, let's look at a hypothetical situation. Let's assume that the operators of a website want to increase traffic to the website. Three reasons why the website might not be attracting more traffic could be:

- The portal's content was not interesting to enough people
- The portal was hard to use, so people stopped trying
- Traffic had not been drawn to the portal effectively

Examining the usage of the site, as discussed above, can reveal how many users visited the site only once or twice. A lack of repeat visitors generally indicates either that users find that the material on the website is not interesting or that the website is hard to use. If, on the other hand, there is a high percentage of repeat visitors but the overall number of visitors is small compared with the eligible population of users, then the problem may be how the website is being promoted.

Let's assume for the sake of the discussion that a high proportion of users, 70 or 80%, visited the site a few times but didn't come back. Two ways of finding out why they didn't come back might be, first, to test the usability of the existing website and, second, to survey users about the appeal of the content. This section discusses the first way, how to carry out user tests to explore whether a lack of usability may be contributing to the problem.

As mentioned above, testing a small number of users will yield very useful information. Nielsen that just five users will usually identify 80% of the usability issues with a site²

2 Jakob Nielsen, "Why You Only Need to Test With 5 Users," Jakob Nielsen's Alertbox, March 19, 2000. (<http://www.useit.com/alertbox/20000319.html>)

The five-user test is qualitative. Its main purpose is to allow the website's operators to look over the users' shoulders and get their perspective about what's working and what's not. Nielsen recommends that for quantitative studies, 20 users are enough. But he strongly advises that the first efforts to assess the usability of a website should be qualitative and gives two reasons: first, quantitative studies are expensive, time-consuming and error-prone and, second, website personnel have to have some experience in user testing before they can perform quantitative studies correctly.

A major problem in testing usability in most websites is finding representative users to test. The ideal situation is when the users are local, so that they can come to your offices to do the tests or, better still, do the tests in their own offices with testers physically present to "look over their shoulders." If neither is possible, user testing can be done using a remote desktop connection and a telephone, so that the tester can see the user's screen and communicate with the user by phone. This is less than ideal, however.

Given the rather special nature of CGIAR centers' websites, it may be possible to define users fairly narrowly, for example scientists and support staff at similar institutions, universities, government agencies, etc., who are interested in the subject matter on centers' websites because of their jobs. It should not be too difficult to find five such people in a city where a CGIAR center has its headquarters. The center's marketing group may have a list of contacts. Nielsen points out several strategies for recruiting test subjects,³ including remembering that there are professional consultants who specialize in just this.

The next thing is to decide what to ask. This depends on the objective. As a concrete example, let's look at what could be asked about the usability

³ Jakob Nielsen, "Recruiting Test Participants for Usability Studies," Jakob Nielsen's Alertbox, January 20, 2003. (<http://www.useit.com/alertbox/20030120.html>)

of the CGXchange website, a site to which everybody in the CGIAR has access. A set of tasks for testing the usability of this website might include navigating to several features that the developers saw as important when they implemented it. The tasks should also range in difficulty from simple to complex.

Here are some tasks that users could be asked to do while you observe:

- Find a job announcement
- Find the CGIAR guide on sexual harassment in the workplace
- Change their password
- Find an article on women and the 2006 East Asian tsunami in the CGVlibrary
- Find the latest issue of the ICT-KM Newsletter
- Find information on how to set up and use a Wiki
- Find something personally interesting to the user

Test sessions usually last up to an hour. The procedure is to sit down with the user and to observe what the user does when asked to complete the tasks. It is often recommended that the test should be done in the user's usual setting, such as their office, as mentioned above.

Usability tester Rachel McAlpine offers some specific instructions for testing that correspond with our own opinions of how such tests should be run:

Prepare a script

A script is part of your own quality assurance system: it helps ensure that you follow procedures, and that you are asking each user to do the same task.

Items on the script:

- *Introduce yourself*
- *Reassure, establish rapport: 'Please think aloud, you cannot make mistakes, we are testing the site (not you).'*

- *Clarify: purpose of test, confidentiality issues.*
- *Start test: open-ended questions, 'what if?', 'tell me more', 'yes?', 'uh huh...', 'mm...'*
- *State task or tasks, e.g. 'What might you click on if you wanted to find out about a home loan?'*
- *Allow users to try to accomplish the task in their own way.*
- *End the test: say thank you, reply to previously unanswered questions, give payment or a gift if appropriate.*

Plan how you will take notes

- *What is your preferred style of note taking? What would make it easier?*
 - *a column for comments on your script?*
 - *blank paper for a mindmap?*
 - *printouts of key pages?*
 - *thumbnails of key pages?*

To save time you can use abbreviations when noting occasions when the user hesitates, looks worried, misunderstands, looks frustrated or gives up.

Test the script

It's an excellent idea to test your script with a colleague acting as the user.

One problem with usability testing is that outsiders may be too polite and eager to please. You can get useful feedback about your own facilitation style from a colleague.

After testing, ask your colleague about the experience. For example, did the colleague feel hustled or hurried or controlled by you? Did the test run smoothly? Did you set tasks in the right order, or the right way? How could you improve your manner?

Now start the test

Sit the user down and run through the first part of your script. Then turn on the computer and show the first web page you're testing.

Start the test.

Notice the user's behavior, and note every occasion the user:

- *hesitates, worries, or is forced to think*
- *misunderstands something*
- *gets frustrated or annoyed*
- *gives up.*

Facilitation style: bite your lip

You have an agenda, certainly: but you are only an observer. Watch the user do what comes naturally. Don't help.

The whole point of the test is to see what users do alone, without you helping.

Much of the time you'll just be probing, and encouraging the user to say what they're doing and thinking.

Use everyday language, not in-house jargon, unless you are testing an intranet.

Keep calm: you want the user to find faults! Don't take it personally.

Suspend judgment: don't even think about solving the problems at this stage.

Keep encouraging the user to think aloud.

Repeat: don't help! Don't think for the user. Explain that you can't answer questions during the test, because it's a test of how people use the site without you. But you'll answer them later if you can. Enjoy the job.

Write everything down. If a task is part of the test, note the time it takes each user.

Don't interfere or ask leading questions.

Check your original task list as you go.

Report on test

Write a 1–2 page report simply noting each problem the user found. Do it immediately while the test is fresh in your mind.

Optional: mark each problem as serious, less serious, or preference. ('Preference' means matters of opinion, such as whether the colours are pretty. Preferences don't affect usability.)

Recommend solutions only if you are required to do so.

Clearly differentiate your recommendations from your observations.

Give your report to the appropriate person or group: e.g. the owner or web development team. Don't delay.

Decide on action

Now you may need to meet with the web development team. Discuss the problems, and decide which ones are easy and cheap to fix, and which need further investigation.

Carry on testing

You'll get better if you keep practicing! Remember, for cost-effective results, test little and often. But just because you are experienced, don't be tempted to speed up the test or put words into the users' mouths. They are the experts on their own experience: not you.

Source: Rachel McAlpine, "Web site usability testing: recommended procedures," Quality Web Content <words that work>, 2006. (http://www.webpagecontent.com/arc_archive/124/5/)

Carry out the procedures above and you will have completed a fast,

inexpensive, but probably very helpful, user study. Finally, read! Educate yourself. Look at the web pages and read the articles mentioned above. User testing is still a bit of an art, particularly because sensitivity to the user's perception, not your own, is paramount. So, immerse yourself in the task and your website's particular characteristics and issues. Then you will get good results.

Also, remember that there are people out there who do usability testing for a living, so you can hire a consultant to do the user testing, if you prefer.

Polls and Surveys

Both usability and usefulness can be measured by user polls and surveys, either solicited on the website or by some other means, such as email. Survey information allows the Web administrator to move from quantitative to qualitative data. "This is the difference between knowing that 8,000 people visited your home page on Monday and knowing that four out of five visitors thought your home page was slow to load."⁴ Internet-based surveys are discussed below.

Keep in mind, however, that usability experts stress that survey research is never a complete substitute for user testing. The reason is simple; verbal responses from users about what they do on websites often do not correspond well with what they actually do. As Jakob Nielsen argues:

Too frequently, I hear about companies basing their designs on user input obtained through misguided methods. A typical example? Create a few alternative designs, show them to a group of users, and ask which one they prefer. Wrong. If the users have not actually tried to use the designs, they'll base their comments on surface features. Such input often contrasts strongly with feedback based on real use.

⁴ Eric T. Peterson, *Web Analytics Demystified: A Marketer's Guide to understanding How Your Web Site Affects Your Business*. Portland: Celilo Group Media and CafePress, 2004. p.10

For example: A spinning logo might look pretty cool if you don't need to accomplish anything on the page. Another example is the drop-down menu. Users always love the idea: finally a standard user interface widget that they understand and that stays the same on every page. However, while they offer users a sense of power over the design, drop-down menus often have low usability and either confuse users or lead them to unintended parts of the site.

To discover which designs work best, watch users as they attempt to perform tasks with the user interface. This method is so simple that many people overlook it, assuming that there must be something more to usability testing. Of course, there are many ways to watch and many tricks to running an optimal user test or field study. But ultimately, the way to get user data boils down to the basic rules of usability:

- *Watch what people actually do.*
- *Do not believe what people say they do.*
- *Definitely don't believe what people predict they may do in the future.*

Say, for example, that 50% of survey respondents claim they would buy more from e-commerce sites that offer 3D product views. Does this mean you should rush to implement 3D on your site? No. It means that 3D sounds cool. The world is littered with failed businesses that banked on people's attitude toward hypothetical products and services. In speculative surveys, people are simply guessing how they might act or which features they'll like; it doesn't mean they'll actually use or like them in real life.

Source: Jakob Nielsen, "First Rule of Usability? Don't Listen to Users," Jakob Nielsen's Alertbox, August 5, 2001. (<http://www.useit.com/alertbox/20010805.html>)

But, of course, surveys are appropriate for measuring usefulness because it is simply not possible to find out what interests a large enough proportion of a site's audience just by observing a few users' behavior.

3. Measuring Usefulness

The Theory

What do we mean when we say a site is useful, and why do we define it as a separate category of measurement? A good start is to ask, Useful to whom?

Tsakonas and Papatheodorou give a good definition of the usefulness of an electronic information system in a recent article:

“Usefulness is the degree to which a specific information item will serve the information needs of the user.”¹

Criteria of usefulness can include such attributes as relevance, format, reliability and timeliness. Fundamentally, usefulness is about the interaction between a user and the content of a website, while usability is about the interaction between the user and a website’s user interface.

Defining a website’s content involves asking questions such as, Who is my website for? What do they want from the website? What do I want from the website? and What can I realistically provide and maintain? Clearly, the designers of the CGIAR centers’ websites have asked and continue to ask these basic questions. The point is that on both sides of the content–user interaction, there are aspects that are specific to each website and its users. Thus, any measurement of usefulness has to take these specific characteristics into account.

It is reasonable to assume that the designers of a website will define their audience and the content of the website, and be able to say why they believe the content is appropriate to the audience. The purpose of

¹ Tsakonas, Giannis, and Christos Papatheodorou, “Analysing and evaluating usefulness and usability in electronic information services.” *Journal of Information Science*, 32 (5) 2006, p. 402.

measuring usefulness is to test whether or not these definitions and beliefs are correct. Most often, the tests involve some kind of survey research because, since the user's expectations and judgments are an intrinsic part of usefulness, and since users differ, asking users is the most convenient way to find out their expectations and opinions. Survey research does have its drawbacks, however, and can be supplemented by less obtrusive measures, such as drawing implications from usage, as discussed above. But, at some point, some kind of user survey is needed.

Types of User Surveys

Internet technology provides several ways to ask users about a website's usefulness. It is relatively easy to create a survey on a web page and to collect the survey data in a database or a spreadsheet. Several online services allow researchers to design survey forms online using only a browser. The survey Web pages reside on the application service provider's site. The application service provider runs the application behind the form, and collects and analyzes the results. Users visiting the website that is being evaluated will find a link that will take them to the survey and run them through it. Finding vendors of these services is easy. Just enter "website evaluation survey" into your favorite search engine.

The link to a user survey may be on a page on the website or in a pop-up box that appears with some predetermined regularity. Website operators can decide whether they wish merely to ask users to take part in such a survey, or whether they wish to offer some kind of incentive. Incentives may be relevant to the site content, such as a publication, or unrelated to the site content, such as a chance to win a gift certificate or other item of value.

Site registration procedures, for example where users supply information in exchange for access to a password-protected part of the site, to a newsletter subscription, or to some other feature, can be used to generate mailing lists of users who can be sent a link to the online survey by email. The link can also be sent to people who may or may not be current users of the site but who **should** be users. For example, a questionnaire about

usage of the main CGIAR website could be sent to a random sample of people who have cgiar.org email addresses. It is also possible to design a survey to evaluate a specific Web page. Often, survey forms to evaluate a single page are built into the page itself and ask very few questions, perhaps only one, for example “Was this information useful?”

Most surveys of website usefulness include questions such as:

- How do you rate the usefulness of this website?
- Did the website provide you with the information you needed?
- If not, what information would you like to see?
- How do you rate the ease of use of the website?
- How frequently do you visit the website?
- Please select the three sections you visit most often.
- How do you rate the content/information available on the website?
- How do you rate the overall design of the website?
- What content or features would you like to see added to the website?
- What functionality (e.g. RSS, discussions, news feeds) would you like to see added to the website?
- How did you hear about the website?
- Overall, how satisfied were you with the website?

Some examples of surveys of website usefulness can be found on the following websites:

- <http://www.cpsc.gov/survey05.pdf>
- http://www.usm.edu/ir/website_evaluation_survey3.htm
- <http://www.cccu.org/about/contentID.28,childContentID.42/about.asp>

As these examples show, each survey can be customized to address questions specific to the content and the perceived purpose of the site. It is also useful to ask respondents questions about themselves, such as their occupation, title and relevant interests. If the survey is large enough, results can be cross-tabulated to show differences among different types of user.

Putting it into practice

A Hypothetical Usefulness Survey

The process of surveying users about a site's usefulness has four phases:

1. Identifying the survey subjects
2. Drafting the questionnaire
3. Administering the survey
4. Analyzing and interpreting the results

Let's look at each of these phases in turn.

Identifying the survey subjects

Carrying out a rigorous scientific survey of users of a public website is extremely difficult, particularly if the potential audience defined for the site represents a very small percentage of the general population. This is very much the case with the CGIAR's public websites. Screening potential test subjects by asking whether they've ever used such a website would generate so many negative responses that it would be prohibitively expensive.

Somehow the population to be surveyed has to be narrowed down, but every method for doing this will raise questions about the validity of the results. However, this can't be helped; you are just going to have to do the best you can and avoid claiming precise levels of confidence or confidence intervals. For a public website that doesn't have an automated procedure such as registration to collect information about people who use the site, one good approach is to find some other list of likely users, such as donors, subscribers to journals, or authors of emails to the site. Another approach is to define an audience very narrowly. In the case of WorldFish, for example, the audience might be defined as members of a global society of ichthyologists.

Generating mailing lists is one of the main reasons for asking website

users to sign up in some way. Such mailing lists are often used for promotional purposes. Users may provide their names and email addresses (or even more information) in exchange for benefits such as:

- A white paper or other specific publication
- A subscription to the site's newsletter
- An RSS feed
- Access to a protected area of the site

Mailing lists generated in this way identify site users, since obviously they must have used the site to sign up. There is a methodological argument against using such subjects because they are 'self selected' or, in other words, they chose to apply for the benefit. But, given the difficulties of finding better subjects, it is probably wise just to accept this limitation and go ahead. You can simply describe the population that you surveyed and make it clear that the population is made up of people who signed up on the site and therefore probably represent the more interested proportion of the total population of users accessing the site.

If the mailing list is large, the number of survey subjects you need to select will depend on how representative and specific the results need to be of the people on that list. Say, for example, you have a list of 600 users who provided their names and addresses in order to receive a monthly email newsletter. To generate results at the 95% confidence level with a plus or minus 5% confidence interval, you would need completed questionnaires from about 240 users randomly selected from that list. If a plus or minus 9% confidence interval at the 95% confidence level is acceptable, you can reduce the sample size to about 100. At the end of the day, however, you will still have to confront the question of how representative the people on the list are of site users in general. If you are simply looking for heuristic, general guidance about the site, you may be able to use a smaller sample size.

To generate a random sample you could either use a table of random numbers or an automated random number generator (for perfectionists, the tables are slightly more random). Once you have a set of random numbers for a population of 600, you can simply list all the website users

on a spreadsheet and select 100, 240, or however many you decide, by the spreadsheet line numbers corresponding to the random numbers.

The most challenging part of the process will be getting your selected subjects to answer the questionnaire. Here are some possible methods to improve the response rate:

- Offer them an incentive for completing the questionnaire (for example, a gift certificate for Amazon.com)
- If they do not respond to the first email requesting their participation, contact each subject again two more times.
- Write a very short, persuasive cover letter in the email

If, after three tries, a subject hasn't answered or has refused to complete the survey, replace this subject with another randomly selected subject until you have the required number of responses. If this requires too much effort, report the response rate with your results and indicate that your results are, therefore, somewhat less reliable.

Drafting the questionnaire

The examples of questions given in the section on Types of User Surveys are fairly good examples of the types of questions you should ask. They are worded neutrally (e.g. not "How great did you think the site was?") and are a combination of open and closed questions.

As a concrete example, however, let's look at a model questionnaire the author created for CGXchange. This, of course, is just an example. The questionnaire for your site will be determined by what you want to ask and your particular circumstances. If you look at the CGXchange website, you may be able to figure out why these particular questions were asked. In particular, the questions ask about specific areas that can be accessed, usually from the CGXchange home page (at the time of writing).

Model CGXchange usefulness questionnaire

1. How did you hear about CGXchange? (Check as many as apply.)

- a. Newsletter _____ b. Email _____ c. Friend or colleague ____
 d. Web search _____ e. Other website _____ f. Mass media _____

2. How do you usually access the Internet (check as many as apply)?

- a. Dial-up connection _____ b. Wireless broadband _____
 c. Wired broadband _____ d. Office Local Area Network _____

3. In your most recent visit to CGXchange, which of the following areas did you access? (Check as many as apply.)

- a. The community of a specific center (e.g. CIAT) _____
 b. The CGX Survey community _____
 c. A collaboration tool (e.g. Wiki) _____
 d. The CGVlibrary _____
 e. News (any kind) _____
 f. Your account information _____
 g. The page to change your password _____
 h. "Connect" information (e.g. staff list, Xperts) _____
 i. CG On-line Learning Resources _____
 j. Knowledge Sharing _____
 k. CG-PM Projects _____
 l. Events (CGIAR consolidated calendar) _____
 m. Policies & Procedures _____
 n. Jobs _____
 o. FAQs _____
 p. ICT Briefcase _____
 q. Other (specify) _____

4. In your most recent visit to CGXchange, how satisfied were you about

finding what you were looking for?

- a. Very satisfied _____
- b. Satisfied _____
- c. Dissatisfied _____
- d. Very dissatisfied _____

5. How easy was it for you to find what you wanted on CGXchange?

- a. Very easy _____
- b. Easy _____
- c. Difficult _____
- d. Very difficult _____

6. In your most recent visit to CGXchange, how satisfied were you with the usefulness of the information you found?

- a. Very satisfied _____
- b. Satisfied _____
- c. Dissatisfied _____
- d. Very dissatisfied _____

7. Overall, how satisfied were you with CGXchange?

- a. Very satisfied _____
- b. Satisfied _____
- c. Dissatisfied _____
- d. Very dissatisfied _____

8. List any content, or kinds of content, you would like to see added to CGXchange.

9. How would you like the CGXchange staff to improve CGXchange?

Not only will your questionnaire differ from this model, but you may also want to add other questions. After all, if you're going to the trouble of sending out a survey, you may want to add a few market-research questions, such as asking the respondents how interested they would be in new features you're considering adding to the website. You don't want to overdo this, however, because your response rate will suffer if your questionnaire is too long.

In the model above, the subjects have not been asked questions about themselves, such as their age, sex, center, job title, etc., because the CGXchange user profile feature already stores this kind of information. In the case of a public website, however, you should add demographic questions that are relevant to your audience, such as their job and title, and perhaps education, gender and age. Generally, subjects are less willing to provide information about themselves than about their opinions, which is why demographic questions are usually put at the end of the questionnaire.

Administering the survey

Once you have a list of people to be surveyed, you have two alternatives for administering the questionnaire. The first is to include it in or attach it to an email. The advantage of including the questionnaire in an email or sending it as an email attachment is that, if it is short, the user may answer right away and simply hit 'reply' to send it back. But there are disadvantages. It may be difficult to format the questionnaire in the email and, more importantly, you will have to enter the results manually into another tool in order to analyze the data. If there are only a few users and just a few questions this may not be too onerous. However, by manually transferring data you risk making errors in data entry.

The other alternative is to use online survey software or an online survey service. Many companies offer online survey services. Basically, these provide a browser-based tool to create a questionnaire, which is then hosted on the service provider's website. You send survey subjects emails that contain a link to the questionnaire, either through the provider's email service or using your own email. Users follow the link and fill out the questionnaire using their browsers. Software on the provider's site analyzes the data and presents the results, which you access from your browser. There is usually also a way to export the data, and the analysis, from the provider's site to your own system. Some companies, such as WebSurveyor, also offer software that can be installed on your website. For the occasional survey though, the online versions are simpler.

Some examples of these online services are:

- WebSurveyor: <http://websurveyor.com>
- QuestionPro: <http://www.questionpro.com>
- Zoomerang: <http://info.zoomerang.com/index.htm>
- SurveyMonkey: <http://www.surveymonkey.com>

As mentioned above, the hard part is getting people to respond to the questionnaire, and we've already mentioned some approaches to that

challenge. However, most CGIAR centers are very likely to have experts in survey research who could be asked to help.

Analyzing and interpreting the results

Online survey services automatically tabulate results. To cross-tabulate results, either among questions or with data from the user profiles, you can use Excel. For example, I would recommend cross-tabulating the responses to questions on access to specific information (question 3 in the model questionnaire) with the responses to questions on satisfaction and ease of use (questions 4 and 5).

I wouldn't be too concerned about precise results, simply because what you're really looking for are ideas to improve the site, not final judgments. Remember that whatever kind of data you're examining, be it usage, usability or usefulness, it won't make any difference unless it's part of the overall creative process of improving the website.

Appendix A: Problems and Limitations of Web Log Data

The Web can be considered mostly ‘session-less’, in other words, each transaction is reported and recorded separately in a log file. When a user types in a URL for a website, or clicks on a link, each request for the file, or files associated with the link, is treated as a single event and is not associated with any other requests issued for files on the site. **Requests for files, not the users, are tracked¹.**

The log recognizes visits by the addresses of machines (IP addresses), not specific people. The number of visitors reported is affected (increased or decreased) by the use of, for example dynamic Internet Protocol (IP) addresses, proxy servers, asynchronous requests made by Flash or JavaScript scripts, caching and download accelerators.

In the case of dynamic IP addresses, a single user at a single machine might actually be using more than one IP address. Alternatively, in the case of proxy servers, all Internet traffic in an organization might be channeled through a server used as a stand-in IP address (often the case with ‘firewalls’ and other corporate computing security arrangements). In this case, all the hits to the website from everyone in that organization could show the same IP address. For example, staff members could access the Internet through a proxy and an Internet node. The files requested by the users—images, web pages and animations—are often cached by the client’s browser. The cache is a temporary storage file. If the user returns to that file (page), his or her computer might retrieve it from the local cache rather than from the website, in which case the server log file would not record the transaction.

The number of ‘hits’ generally reported is the number of files requested, which may or may not correspond to pages (a ‘page’ on the site might

¹ <http://depts.washington.edu/pettt/papers/arthritis/Ramey.doc>

for instance contain several graphic images stored in separate files, each of which is logged separately). Thus, reaching conclusions about users and what they are doing requires making complicated logical links and inferences that are far beyond what the actual data at hand shows.

The information provided by web servers is incomplete. Web servers only record requests to the server, so any request that does not reach the server will not be recorded. This mostly includes requests for files cached by the user's ISP, company proxy server, or browser. It's also not possible to determine click-stream data from web logs, such as the amount of time spent on each page. IP addresses give an impression of which types of organizations are accessing a website, but they are effectively useless for identifying individual users.

HTTP (hypertext transfer protocol) is an anonymous and stateless protocol² which means that users are not individually identifiable, and that each request is independent and not explicitly connected to other requests by the same user (unless the website requires a password login, for example, or uses some other method of state-management, such as cookies).

The server does identify the IP address of the computer making the request (so that it knows where to send the response), but this IP address does not necessarily relate to a single unique user, as explained before. This leads to over- and under-counting.

Finding a way out of this predicament with any degree of certainty is difficult because it is not possible to measure how many unique users accessed a website. It also means that it is not possible to deduce how many 'visits' were made to a website. Clearly, as it is not possible to determine unique users from the IP address alone, attaching requests to an IP address and compiling them into a single session is also a meaningless activity. Cookies are often cited as a realistic way to get around the problem of identifying users, but it is not possible to rely on them to produce valid reports.

² <http://library.n0i.net/programming/web%20programming/rfc2616/>

Visits vary in length; sometimes users will not request a page for a long time, but will still consider themselves to be visiting the site. They may even be paying close attention to the site all along, for example they may be reading a long article in detail. This, plus the issue highlighted above—that not all requests actually get logged in the first place—makes it clear that allocating a meaningful ‘timeout’ period for the purpose of defining visits is almost impossible.

Log data includes statistical noise. Basic log file analysis must identify and remove traffic that artificially inflates the ‘real’ figures. This includes requests made by robots (such as Google), and requests that were unsuccessful (those returning a ‘404’, for example). It might also be desirable to filter the analysis of actual files being requested, for example by ignoring requests for image files and style sheets.

A server can only record the requests that it receives. When a user clicks on a link to request a page from a website, that request may never reach the server, as the files requested might be cached by the browser, the user’s ISP or company proxy server. This has a number of effects. The log will be undercounting the ‘real’ number of page impressions that a site is generating, and this will generally affect the count on the more popular pages on the website, as these are the ones most likely to be cached.

Because web logs are incomplete, making assumptions about the user’s actual route through the site is highly problematic (even assuming that we can identify individual users). Current versions of browsers can open many web pages simultaneously in multiple tabs in the same window; the web server will record each request sequentially, making the ‘path analysis’ metric nonsensical. It is common practice to use the back button to navigate around a site, viewing a page, and then going back a step, perhaps to an index, and then viewing another page. By using the back button in this way—incidentally an excellent usability tool—popular pages, and in particular ‘index’ pages which contain navigational links, will often be undercounted, thereby presenting us with a false ‘click-

stream'. Nevertheless, much analysis of log files incorrectly assumes that they contain a complete picture of user activity.

It is possible to design website pages that instruct client machines not to cache them, and this is sometimes advocated as a potential solution to the problem of incomplete log data. This is, however, probably not advisable. Caching improves the user experience by speeding up download times, and reducing the workload of the web server. While there are legitimate reasons why a designer might wish a web page not to be cached, for example the homepage of a frequently updated news site, collecting traffic data is probably not one of them and, in any case, due to sophisticated caching techniques by ISPs, the data collected would still most likely be partly incomplete. It is better to know that the data is incomplete and handle it accordingly, than to incorrectly assume that it is authoritative.

Individual user information is misleading. As HTTP allows users to access files anonymously, the server will only record the IP address of the machine making the request. In the case of many ISPs or company proxy servers, this IP address can represent large numbers of individual users. Log files will not show how long a user has spent on a website, or how long they spent viewing each page. Certainly, we know the time of each request logged but, as we now realize, this is an incomplete picture of their route through the site. Even if it were possible to assume that the log recorded all requests, it cannot record the time at which the user stopped looking at the last page, or left the site for another.

The server is unable to log where the user goes after leaving the site as, by definition, that request is sent to a different server. Time spent on the last page of the visit is unknown, and therefore the total time spent on the site cannot be determined. The possibility that the last page viewed might well be the one that held the attention of the user the longest can only add to this uncertainty. Certain sites, such as online banks, enforce a session log-in/off, but this does not tell us how long the final page was truly viewed, as commonly users will allow the last page to time out, or simply close down the browser window. In any case, this solution is not appropriate for most sites.

If we define the boundaries of what we know, then we can also define what we do not know, and the information becomes all the more valuable as a result. The greater the distance between the actual data and the conclusions that you draw from them, the less reliable and certain the conclusions are, and the more caution you need to exercise in making design decisions based on them.

For all these reasons, log file analysis is not an efficient way to improve the usability of a website.

Keeping in mind these challenges to data interpretation, it is important to have a brief look at the data reported in server logs.

Log file formats

Log files follow one of two formats: common log file (CLF) format and extended log file (ELF) format.

The **common log file** (CLF) format basically records the date and time of the transaction, the IP address of the visitor, the size of the file in bytes, and the status of the request (for example, '404, file not found').

The **extended log file** (ELF) format³ includes other data about visits and visitors, such as the visitor's browser and platform, the website that the visitor is coming from (called the 'referring page'), or the search term the visitor used, if it was entered from an engine or directory. The ELF format can include 'cookie'⁴ information as well, if available.

³ <http://www.w3.org/TR/WD-logfile.html>

⁴ A 'cookie' is a small file that records a visiting computer's activity on a site. When a user visits a site, the site server can transmit a cookie file that records the files that the user has requested. The cookie is placed in a folder on the user's computer and, when the user re-visits the site, the server requests the cookie file that it sent to user's machine before and updates it with data about current visit. In this way, the server can build up a historical record of a computer's actions on the site.

Appendix B: How to Use Google Analytics

To use Google Analytics you must first register.

For this use the email address that is associated with the domain you want to track.

After this you will be able to access the Google Analytics Web application. This application gives you the option of registering different domains for the service. Follow the steps to register the domain you want to track. The result will be a fragment of JavaScript code. This must be added to every site in your homepage between the HEAD and the BODY tag. Ask the developer of your site to do this for you.

After 24 hours you will be able to see the first results from Google Analytics for the tracked domain.

These will be updated every 24 hours. With Google Analytics you are not able to track in real time but for most sites real-time tracking is not necessary.

Google Analytics can give you a lot of interesting information, for example:

Standard Values

Standard values include:

- How many (absolute/unique) users visited your site?
- Average pages per visitor.
- Average time per site.

These values are presented in simple graphs. You will see them on the main dashboard. But these values are for the entire website and thus are only useful for the big picture. To get more details you have to filter this data for different sections of your site.

Languages used from the user browser

This is important information if you want to optimize your page for different languages. Not only will it help you to decide how many languages you want to support but, if you find you have a lot of users for a specific language, you can use this information to optimize your text for this group, or find better ways to present information to them.

Browser, operating system and other technical information

Google Analytics gives you a lot of technical information, but how can you make use of this? At the beginning of every Web-based project the developers look at optimizing the application for a particular environment. They list technical issues such as screen resolution and the browsers that will be supported.

For example:

- You define a low screen resolution as the default but a typical user is using a much higher resolution. Your site will be classified as 'old fashioned and not state of the art'.
- You use too many pages to present a limited amount of information so the probability that a user will leave your site is high.
- The website must be flexible and easy to change.

Google Analytics provides information on the operating systems and browsers that are being used by visitors to your website. You need this information when making modifications to your website.

Every time you make modifications it is a good idea to do a compatibility test with W3C tools because every combination of browser and operating system works a little differently. These tools are available at www.W3C.org. The only way to be sure that the target group is able to use all features in a website is to run tests in the same environment that they are using. Normally it is not too expensive to create test environment combinations for different hardware platforms.

These are called virtual test environments. To build these environments it is necessary to equip a few PCs with virtualization software, for example VMWare. With these it is possible to create different combinations of operating systems and browsers.

Google Analytics checks two other important sets of client information.

The first is the version of Flash that is installed inside the browser. Flash is a technology for creating dynamic websites and is often used by design-oriented website developers. But there are some critical factors to think about when using Flash. Flash has many different versions and websites created with Flash are not easy to maintain in several different environments. Another issue is scalability. Performance is limited, tracking user actions is expensive and significant design changes are time consuming.

The second important piece of client information is the version of Java. Java is a hardware-independent programming language and comes in many different hardware–software combinations. If users have Java it is possible to build complex Web applications, such as enterprise environments, for them.

Infrastructure information

Infrastructure information includes:

- Bandwidth
- Networks and backbones

The information about the infrastructure the user is using to access the Internet is important for the technical design and ‘feel’ of your website. Bandwidth is expensive. This means, for example, that in Europe bandwidth is less limited than in Africa. So, if the main user group of a website is in Africa the website must be optimized for low bandwidths. Optimization is independent of website content.

The entire website does not have to be optimized for every user group. You can identify the sections of your website that are relevant to each group and optimize them appropriately, or you can create different versions for different bandwidths. The switch between both versions must be easy, maybe by a button on every Web page.

Keywords and why they are important

Everybody uses keywords to search for information with Google or other search engines. Information about keywords is useful in generating content that can be found with these keywords. This sounds trivial but hardly anyone actively uses this information.

Google Analytics will create a list of keywords that increases hits on the tracked website. A top 100 list of keywords can guide rewriting the content of a website.

Keywords change over time so it's important to track keywords for the most important Web pages to make sure that they hold their ranking inside the search engines. To find the keywords relevant to your website use a thesaurus and check other sites that may also be associated with these keywords.

For the most important keywords it is useful to build up test pages and monitor the traffic. These are the first steps towards increasing a website's ranking in the search engines. More information about this can be found by searching for 'SEO' (search engine optimization). Take care with these techniques as some are very sensitive and incorrect handling will be punished by strong negative rankings in Google searches.

Geographical information

Google Analytics can provide a lot of geographical information, for example from which country and city a request was made. This is a useful tool to track whether local events are changing the traffic. Normally, a few days before, during and a few days after a conference, the traffic

from a city where a conference is being held is a little higher. The same happens with other events. This is a good indicator of the effectiveness of advertisements, announcements and other regional promotion. Each region can be tracked independently and compared with other regions.

Appendix C: Website Performance Test

Every web application must be tested before being officially launched. These tests are important in making sure that the application is working in the way you expect.

One part is the user test. For example, Is the user able to find the information he or she is searching for? The other part is technical. For example, What will happen if more than one person is using the web application? Is the server able to handle 1000 simultaneous users? How fast is the application reacting? To answer these questions, you must build up test scenarios. What actions in which order are you expecting from a standard user?

Tools to do such tests are available. One is called WAPT and is very easy to use. The idea is very simple. You use your web application and WAPT records your actions.

In this way you can build up a range of 'standard users'. Then you can use WAPT to count different users working at the same time, increase and decrease of the number of users and so on. The documentation of this tool is good and easy to read for non-technical persons.

After each test you will get a detailed graphical report with information about performance, errors, the volume of transferred data and so on.

The advantages are:

- The tests can be done after every change in the Web application without involving real test users.
- The results are comparable because the tests are identical.
- Changes inside specific workflow processes are easy identified.
- The tests can be done from different geographic locations.

Step-by-step it is possible to build up a test pool for the complete Web application. This is a good way to increase the quality of the Web application from the business/non-technical point of view.

Appendix D: Glossary of Terms

Log file analyzers¹ count Web log traffic by hits, page views and visitors (users). Each is counted independently of the others, and each has its own advantages for analyzing the traffic.

Hits. Hits are accepted log entries. So, if there are 5,000 entries in the log file and there are no log filters, and all the entries are valid (i.e. none of them have corrupt dates), then the application will report 5,000 hits for the file. If there are log filters that reject certain log entries, then those will not appear as hits. Log entries that are accepted by the log filters will count toward the hit totals. Because there are no default filters that reject hits, you will generally have nearly as many reported hits as you have log entries.

Page views. Page views correspond to hits on pages. For instance, if you're analyzing a Web log, and a hit on /index.html is followed by 100 hits on 100 images, style sheets, and JavaScript files that appear on that page, then all these will count as a single page view—the secondary files do not add to the total. This is implemented in the log filters—page views are defined as log entries that are accepted by the log filters, and that have a page-view value set to 1 by the log filters. Log entries that are accepted by the filters but that have a page view of 0 set by the filters do not contribute to the page view total. Therefore, you have complete control over which files are 'real' page views and which are not—if the application's default filters do not capture the preferred definition of page views, you can edit them until they do. By default, page views are all hits that are not GIF, JPEG, PNG, CCS, JS and a few other types of file. See Hits, above, for more information on log filters.

¹ Source: Sawmill's web site

Visitors or Users. Visitors correspond roughly to the total number of people who visited the site. If a single person visits the site and looks at 100 pages, this is counted as 100 page views, but only one visitor. By default, some applications define visitors to be ‘unique hosts’—a hit is assumed to come from a different visitor if it comes from a different host name. This can be inaccurate due to the effects of Web caches and proxies. Some servers can track visitors using cookies if the Web logs are set up to contain this information.

Sessions. Sessions are similar to visitors, except that they can ‘time out’. When a visitor visits a site, and then leaves and comes back later, it will count as two sessions, even though it’s only one visitor. To reduce the effect of caches that look like very long sessions, log file applications also discard sessions longer than a specified time.

Session user.² This is the session of activity that a user with a unique IP address spends on a website during a specified period of time. The number of user sessions on a site is used to measure the amount of traffic a website gets. The site administrator determines what the time frame of a user session will be (e.g. 30 minutes). If the visitor comes back to the site within that period, it is still considered as one user session because any number of visits within those 30 minutes will only count as one session. If the visitor returns to the site after the allotted period has expired, say an hour from the initial visit, then it is counted as a separate user session.

Other ways that site administrators measure the amount of traffic a website gets are by unique visitor, hit, click-through and page view.

Referrers. URLs denoting the portals or websites through which another site is reached.

Unique Users. Individuals, often identified through the use of cookies, IP addresses, or passwords, who visit a site.

² http://www.webopedia.com/TERM/U/user_session.html

EVALUATING YOUR WEBSITE

A guide for CGIAR Centers to evaluate the usage, usability and usefulness of their websites

Measuring the 3Us (usage, usability and usefulness) of your website is key to making sure that you are meeting the objectives and impact you set out to achieve when you built your website. Knowing how many people visit your site, who they are and what they do while they are there (usage) will help you tailor your site to deliver, share or pull in the information or messages your audiences most need, in the way audiences want to receive and contribute to it. Knowing how easily visitors find what they are looking for and their perception of your site (usability) will help you improve its functionality and the user-experience—encouraging more use of your site. And knowing how well your site meets your visitors' information needs (usefulness) will help you improve both your content and its organization to meet those needs.

This guide provides an introduction to measuring these 3Us together with suggestions about how you can use Web analytics, surveys and other instruments to improve the 3Us and ultimately the impact of your website. Examples provide a step-by-step guide to putting this knowledge into practice.

This document is targeted at the broad range of people who contribute to the CGIAR's websites, including communication specialists, information managers, management and IT technical specialists.

Massimo Buonaiuto



Collaborate, Create, Communicate