IRS Seeks to Develop New Web-Based Measurement Indicators for IRS.gov

Diane M. Dixon
Statistics of Income, IRS, P.O. Box 2680, Washington, D.C. 20013-2608
diane.m.dixon@irs.gov
Presented at the 2003 American Statistical Association

Key Words: Website Analysis Tools, Website Activity Assessment

In 1996, the IRS created and implemented its own website—irs.gov—to allow taxpayers easy access to IRS information and resources at their own convenience. Since the site's inception, the IRS has relied on web analytics to assess the site's usefulness and to make improvements to enhance the customers' experiences and satisfaction. Serving customers and improving customer satisfaction within the diverse customer base of the IRS is a difficult task, but one to which the IRS is fully committed.

Although the goal is simply stated, there is no single approach to understanding the successfulness of a website or the level of satisfaction associated with it. With this in mind, the IRS has utilized several tools, including focus groups and customer surveys. However, in order to assess satisfaction on a large scale, the IRS has learned that understanding the underlying web activity is the key to designing a website that meets its customers' needs.

Understanding customers

As the customer demand for more functional websites increases, so does the need to understand how site usage affects an organization. There is a plethora of customer data that can be collected and used to interpret site usage. For some sites, a demographic customer profile is important. Such information can help an organization define its market, which can aid in attracting new customers and generating revenue. However, simply collecting various demographics about customers will not result in a better understanding of customer needs. In order to understand customers, one must analyze customer behavior. Customer behavior data afford web administrators the ability to retain customers and predict future customer relationships [1].

Making the decision to profile customer behavior is the first step; however, to make it work, an organization must first consider its specific needs in order to tailor results that will help with decision making and planning. In order to choose measures that will be valuable, certain questions must be addressed, including, What is the website's purpose? How are website changes decided upon currently? What makes the website successful? Addressing these questions will help narrow down which measures will be most valuable to assess a site [2].

Introduction to web analytics

Following the evolution of technology, the way in which website traffic is analyzed has advanced greatly in the last decade. From primitive measures such as hit counts and files downloaded, web metrics have blossomed into a variety of different tools that are valuable both independently or combined into a suite of analysis tools. Depending on the data collection software, web administrators can collect the number of visits, unique visitors, and page views associated with a site, as well as various other web metrics, including path analysis, referral pages, and an assortment of customer demographics, while still collecting hits and downloads.

Since the launch of irs.gov, the IRS has recognized the importance of monitoring site activity. Using two of the most common web metrics at the time—hits and downloads—the IRS collected data to describe the web traffic on the site. In January 2002, the IRS launched WebTrends Reporting Center®, which gathers raw website data and transforms it into a collection of reports easily accessed via the Internet. WebTrends® has allowed the IRS to capture more site data, providing more insight to customer behavior. These additional metrics include visits and page views—displayed for a day, week, month, quarter, or year, depending on preference and need.

Metric analysis

Hits: A hit is a file that is requested by a visitor's computer. Each individual webpage consists of numerous hits—the HTML page itself counts as one hit, but each graphic or hyperlink is also

interpreted as a single hit. The amount of hits on each page is dependent upon the page design. The intended use of this metric is to measure website server workloads—how much stress is placed on a server due to site usage. Depending on the size of the server, the amount of file requests could have a serious impact on the performance of the server, as well as the availability of the website. Therefore, knowing the volume of hits related to irs.gov is important to IRS information technology (IT) personnel. Using this data, they can assess server performance and make decisions concerning equipment needs.

As previously mentioned, the volume of hits is proportional to the design complexity of the website. Each individual page may consist of a varying number of hits, meaning graphic- or link-rich pages produce higher counts than simple pages which yield lower hit counts.

Due to limitations of the current version of WebTrends® running on the IRS system, hit counts for individual pages are not available. However, if these data were available, one would see the same number of visitors produce a higher amount of hits by visiting the "Where to File, By State" page of irs.gov than they would if they visited the "Retirement Plans—Educational Services Program" page [3], [4]. Both pages include all of the links contained on the top and left navigation bars, but the "Where to File" page has a graphic of the United States that contains 50 links, as well as a listing of each state, adding another 50 links. It also has a few other links and graphics. However, the page about the educational services program only consists of plain text and two other links. The significant difference in the number of links and graphics on the pages will notably alter the number of hits associated with each page, even if both pages are visited an equal number of times.

Downloads: A download occurs when a file is copied from the website server to the user's computer. Files are identified by their file extensions (e.g. .xls is the file extension for a Microsoft Excel® file). Web analysts can program the software to count certain extension types so that they can filter out types of files that they do not want to include in the analysis.

For sites with numerous downloadable files, this metric can be extremely helpful in determining what is important to the majority of customers. This type of analysis can help site designers redesign navigation in order to guide customers to more popular files and products. Downloads can also illustrate the effectiveness of recent marketing campaigns. By using historical data, analysts can calculate the increase in downloads for files promoted in campaigns and then determine the success of the campaigns.

Using download counts, one can also determine which files are accessed often, and which are not. This can help site designers analyze the setup of the current site. Files with the least number of downloads may be expected to be found at the bottom of the list, due to their age; however, if a designer expects more customers to access certain files that are currently not being accessed, the designer can alter the way in which these files appear on the site, to help improve accessibility. Then, using current and historical data, an analyst can determine if this change was helpful to their customers.

This metric also allows analysts to see trends among the types of files downloaded during certain times of the year. Customers may want different information, depending on what month it is. This is certainly true for the IRS—the majority of IRS file downloads are predictable, following the filing seasons. However, there are portions of irs.gov, such as Tax Statistics, that are not as foreseeable. Files contained within Tax Statistics are produced by the Research, Analysis, and Statistics organization within the IRS. The way in which customers access these data files is not predictable. However, analyzing these data over time has allowed the web designer to better understand what customers want and when they want it. This knowledge has led to the discussion of designing navigation based on the time of year—using the landing page of Tax Statistics to spotlight certain data, making it easier for customers to locate desired information.

Visits: A series of actions that begin when a customer lands on his or her first page of the website and ends when s/he either leaves the site or remains idle for more than 30 minutes is considered a visit. The "number of visits" may include multiple visits made by the same user.

In order for a visit to be tracked and counted, it is not necessary for a user to begin on the site's landing page. This is essential for many websites since many customers utilize the bookmark function for pages within sites that they visit often. Certain types of browsing behaviors, including jumping around a site, refreshing pages,

and wrongly selecting pages, can greatly influence certain measures, leading to inaccuracies; however, these behaviors have no affect in the measurement of site visits. This ability makes the number of visits a valuable statistic to most website analysts.

Using this metric, a web analyst can determine how many visits are made to the site within a certain timeframe—an hour, a day, a month, a quarter, or a year. The number of visits can be analyzed historically to determine customer growth. Since it is possible to gather these data based on the time of day, this metric also allows IT personnel to determine the slowest periods of customer activity so that system upgrades and changes can be performed at a time that does not affect a large number of users.

Unique Visitors: Although visits are important in assessing a website, many businesses are interested in how many unique people are visiting their sites. Calculating this number allows a company to further determine the usefulness of its site.

This is a breakdown of the number of visits, allowing one to see how many individuals are behind those numbers. Tracked correctly, one could use this measure to determine the number of customers who visited a website within a certain timeframe. This differs from visits because no matter how many times a customer visits within the timeframe, they are only counted once.

The ability to identify unique visitors also allows web analysts to assess repeat visitors, which further illustrates the usefulness of a site. The measure of repeat visitors may indicate satisfaction among customers, which may reduce or eliminate their need for an alternative site from which to obtain information.

Page Views: Each HTML page is tagged as a page. When a visitor accesses a page, it requests all of the hits on that page, including the page itself. In order to report the number of page views, the website analysis software separates the page hits from the other hits. These numbers make up the page view metric.

Much insight can be pulled from this statistic. One can assess which pages are accessed most, as well as those that are not. Although one cannot assume that the pages with the most views are the most useful to customers, these data can be useful during site redesign. If site owners have a general idea of what information is most appealing to customers, they will be able to determine if visitors are finding that information. Low page views for such pages could be an indicator of site navigational problems. This is similar to the information that downloads provide; yet customers need not download anything to obtain information concerning their interests. This measure equates sites with copious amounts of downloadable files to sites with few or none, thus allowing comparison between these two site types.

As with downloads, page views can also be helpful when determining if customers access types of information at certain times of the year, allowing for further navigational improvement.

Additional Metrics: Although the metrics described above do provide an immense amount of insight into customers' web-browsing behavior, there are other metrics that can further detail website usage, providing a more in-depth understanding of one's customers.

Some software packages allow web analysts to track paths to certain information within a site. By monitoring these paths, analysts can determine if the site navigation is allowing customers to easily access information.

Another valuable tool is one that captures referring pages—the page the customer used to link to a site. Using this feature, one can determine which search engines are most popular among the majority of users. Web managers can then contract with those engines to have their links appear closer to the top of certain searches. This metric also allows analysts to assess the success rate of certain partnerships with other sites, as well as whether or not that partnership should continue.

Although all previously mentioned metrics have focused on customer behavior, software can also collect demographic information about customers. This includes geographic regions, countries, cities, organizations, and domain names. Such information could be useful in various ways. Demographic information can help a site designer tailor a website to the audience. Understanding the audience and designing a site specifically for them will help attract customers and generate first stage revenue [5].

Limitations with web analytics

Depending on website environments, policies, and restrictions, the usefulness of web analytics can be quite limited. Though the data might be insightful, analysts may not be able to fully appraise their sites using certain metrics, even in an unrestricted environment.

Interpreting behavior: Complete interpretation of this data relies on some assumptions, which may not be reliable. For instance, the most downloaded file for a certain timeframe does not indicate that the file was useful to the customer, or even if it was what s/he was searching for. This concept also applies to other metrics, such as page views. Certain pages may be viewed frequently enough to appear in the listing of the top 50 pages viewed; however, this page may not be useful to most customers—it may even be an intermediate page that must be viewed before gaining access to any number of files. (For example, on the landing page of irs.gov, there is a link to the "Where's My Refund" feature. This link takes the visitor to an intermediate page that explains the information necessary to proceed. At the bottom of this page, there is another link that goes to the actual feature.) Using these assumptions, it is possible that a poorly-designed site could produce a significant amount of page views and downloads, which may lead some people to believe that the site is better than one that produces less because its navigation is better.

Another inherent problem is that web-browsing behavior can vary greatly among customers. Experienced Internet users may view less pages, download less files, and spend less time overall on a site. These users may also visit less frequently, as they may find everything they needed in one visit; whereas, inexperienced customers may need to make several visits before finding everything. While web metrics may indicate otherwise, this behavior may not necessarily signify that their satisfaction with the site is lower.

Cookies: A cookie is a small text file placed on a customer's computer hard drive by a web server, usually unnoticed by the customer. This file allows the web server to identify individual computers—enabling a company to recognize returning users, track online purchases, or maintain and serve customized web pages. Cookies can also facilitate the collection of personal information, such as extensive lists of previously visited sites, email addresses, or other information to distinguish individual customers [6]. The Privacy Act of 1974 set regulations concerning the collection of personal information from a citizen [7]. Persistent Internet cookies are considered personally identifiable information and, thus, are covered by this Act. In 2002, the E-Government Act formally delegated responsibility to the Director of the Office of Management and Budget (OMB) to establish government website policies [8]. However, even before the 2002 Act, OMB established a cookie-free policy, explained in Memorandum M-99-18 [9]. In January 2002, the Department of Treasury clarified the policy, explaining that "persistent cookies shall only be granted when the bureau or office has presented documentation which details a compelling need to gather necessary data on the subject website" [10]. The inability to use permanent Internet cookies seriously restricts data interpretation. Without cookies, web analytic software must rely on Internet protocol (IP) addresses in order to collect data about customers. An IP address is a 32-bit numeric address written as four numbers separated by periods. This address is related to an Internet Service Provider's (ISP) server. Large ISPs, such as America Online (AOL) and the Microsoft Network (MSN), have millions of customers sharing numerous servers, meaning that a single IP address may represent thousands of people. For example, if five AOL customers access irs.gov, they may be recognized as one, two, three, four, or five customers.

This notion has a serious affect on website data, especially since most IP addresses are dynamic (temporary) rather than static. This means that the majority of web users have a different IP address every time they visit a site. The problem is made worse by ISPs that allow a client's IP address to change with every new page, meaning that every page view will register as a new visit.

Caching: A cached file is one that has been previously stored on a system (e.g., a personal computer or an ISP server), making reuse of the page or object easier on the customer. When a visitor re-accesses a page or file that has been cached, their system accesses it from the cache location rather than the main web server that hosts the file. The objective of caching is to make efficient use of resources. Although this computer practice may positively affect a customer's experience when accessing a file (e.g., by significantly lessening the download time), it does negatively affect the site's web analytics, as hits, downloads, and page views of cached files will not be captured.

File transfers: As mentioned previously, content- or file-heavy websites greatly rely on data concerning downloads. Such data can provide website owners with the best insight into understanding their customers. Sites with years of historical files, like irs.gov's Tax Statistics, are interested in understanding how downloads change over time, relying heavily on historical web analytics.

The problem with file transfers is that depending on the software package, the way in which files are sent may differ. Some software packages allow all files to be sent as a single file—this is the ideal method of data transfer. With this method, 1,000 downloads correspond to 1,000 actual downloads. However, other software packages split a single file into multiple packets, each registering as an individual download, which greatly inflates the number of downloads reported. With this method, using the example above, 1,000 downloads reported represent the total number of packets sent, which corresponds to a much smaller number of actual files downloaded, depending on how many packets each individual file was split into. The latter method makes interpreting downloads more complex, leaving analysts to rely on other metrics to evaluate their sites.

Educating data users

As explained above, the usefulness of web metrics can be severely restricted. Because of this notion, and a general confusion and lack of education surrounding web metrics, the IRS has begun an effort to educate website managers on definitions and usage of these metrics, as well as how certain limitations impact data interpretation. Only when there is an understanding of the data can it be utilized in such a way as to help improve the site and make more accurate interpretations of customer behavior.

To initiate this learning period, the IRS solicited information from members of the Web Facilitation Group (WFG)—a group of IRS employees responsible for setting IRS website policy—concerning how they use current data, the types of data wanted, how they plan to use additional data, what types of reports they generate using current data, and what types of decisions are made using web statistics. With this knowledge, the IRS will be able to determine the current level of knowledge among the WFG and decide where the education process should begin.

Future discussions with the WFG will focus on how irs.gov web statistics can and cannot be used to interpret customer behavior. Once members of the WFG have a better understanding of irs.gov data, they will be able to provide more accurate reports for their colleagues and ensure that statements about the site are correct.

Developing detailed reports

To aid the web statistics educational process, the IRS plans to develop new reports for irs.gov data. The new reports will contain a significant amount of annotation, allowing for easier and accurate interpretation of data. The IRS plans to develop individual reports for each of the IRS business operating divisions (BODs), as well as a report for all of irs.gov. By including definitions of certain measures, providing an initial data analysis, briefly explaining uses of each measure within these reports, and explaining the impact of limitations, the IRS hopes to help the BODs make well-informed decisions concerning their respective sections of irs.gov, limit the amount of misinterpretation, and distribute the most accurate reflection of website usage.

Upgrading statistical software to improve usability

In conjunction with the education effort, the IRS recently started researching new software options that offer additional functionality, as well as eliminate some of the limitations that currently hinder the interpretation of customer behavior. With the current version of WebTrends®, the IRS cannot generate metrics for individual BODs. Instead, the software produces most data general to the whole site. As one would expect customer behavior to vary in each portion of the site, this makes customer behavior interpretation much more difficult.

Though most of the data generated by WebTrends® is whole-site-specific, the IRS can program the software to gather certain data specific to individual sections of irs.gov; however, this capability is still quite limited. With an upgrade, the IRS will be able to collect web statistics for various sections of irs.gov with ease. This new ability will aid the development of individual reports, as mentioned above.

Conclusion

As one of the most powerful tools used to disseminate information, the Internet has created a

world of faceless customers—people who seek information at their convenience. IRS.gov allows taxpayers 24-hour access to forms and filing information, which reduces the number of calls made to IRS call centers, changing the way in which taxpayers interact with the IRS. However, in order to sustain the success of this type of relationship, the IRS has to recognize the necessity of understanding web customer behavior.

By utilizing a software package to gather data on customer behavior, the IRS has been able to acquire, build, and sustain solid customer relationships without truly interacting with its customers. However, having these numbers alone is not the solution to interpreting customer behavior. IRS web analysts must understand the metrics, as well as the limitations associated with each. Education is a must when distributing reports about web analytics, as without such knowledge, misinterpretation of data is to be expected.

When utilized, analyzed, and interpreted correctly, web analytics can lead to a significant improvement in the usefulness and success of a website, allowing the IRS the potential to attract new customers, retain others, and maintain a high satisfaction rate among all.

Notes and References

[1] Novo, J. (2002), *Drilling Down: Turning Customer Data into Profits with a Spreadsheet*, Bangor: Booklocker.com, Inc., Chapter 1.

[2] WebCriteria, "Executive Guide to Improving Website ROI: Questions Every Business Manager Should Be Asking," WebCriteria Online, available: http://www.webcriteria.com/company/white_papers/index.cfm. [July 11, 2003].

[3] The IRS "Where to File, By State" page can be accessed at *http://www.irs.gov/file/content/0,,id=105693,00.html*.

[4] The IRS "Retirement Plans—Educational Services Program" page can be accessed at *http://www.irs.gov/retirement/article/0,,id=96272,00.html*.

[5] Novo, J. (2002), *Drilling Down: Turning Customer Data into Profits with a Spreadsheet*, Bangor: Booklocker.com, Inc., Chapter 1.

[6] An example of a well known website that uses cookies is Amazon.com. After visiting Amazon for the first time and performing a simple search for one of its many products, Amazon will tailor its main pages to better suit the customer's needs. This tactic is used with the hope of selling customers additional items that they may not have otherwise purchased.

[7] 5 U.S.C. § 552A (1996).

[8] H.R. 2458/S. 803.

[9] Lew, Jacob J. (1999), M-99-18: Privacy Policies on Federal Web Sites [Memorandum], Office of Management and Budget, available: http://www.whitehouse.gov/omb/memoranda/m99-18.html, [July 11, 2003].

[10] The Department of Treasury, (2002), 81-08: Certification Process for the Use of Persistent Cookies on Treasury Web Sites. [Treasury Directive], available: http://www.ustreas.gov/regs/td81-08.htm, [July 11, 2003].