



MERCURY INTERACTIVE

WHITE PAPER

ActiveWatch: Challenging Industry Perceptions of Web Performance

MERCURY INTERACTIVE CORPORATE HEADQUARTERS

1325 Borregas Avenue Sunnyvale, CA 94089
tel: 408-822-5200, 800-837-8911 fax: 408-822-5300

www.mercuryinteractive.com

www.mercuryinteractive.com

ABSTRACT

In the past few years, analysts and industry experts have developed several guidelines for Web application performance, including the 8-second rule, that have been adopted as standards by many organizations. Many of these standards, however, are based on perceptions, not hard data. Nonetheless, they have greatly influenced performance management decisions made by numerous IT and operations groups.

During the past year and a half, engineers at Mercury Interactive have monitored hundreds of Web applications worldwide with ActiveWatch™, a 24x7 application performance management service. The ActiveWatch monitoring data reveals new patterns in the way complex Web infrastructures affect the performance experienced by end users. In addition, it challenges many of the industry's long-standing beliefs about Web performance, including:

- What is considered acceptable application performance—from the end-user point of view
- How availability is experienced by end users
- How long users will wait for pages to download
- How to define performance thresholds
- Why intermittent performance problems occur

Organizations can leverage this ActiveWatch data to launch more effective performance management programs that maximize Web application performance and provide positive end-user experiences.

CONTENTS

Introduction4
Redefining Web Application Performance5
Root-cause Analysis: Diagnosing Problems9
Summary11
About Mercury Interactive11

INTRODUCTION

Mercury Interactive’s ActiveWatch is a managed service that proactively monitors Web applications. By simulating real-world user actions on a continuous basis, ActiveWatch provides a clear picture of a Web application’s end-to-end response times and the impact of performance problems on end users. ActiveWatch also provides customers with advanced warnings of performance problems as well as the ability to quickly analyze their root causes.

As a fully outsourced service, ActiveWatch is managed by a team of experts at Mercury Interactive’s network operations center. During ActiveWatch’s first year and a half of operation, these experts collected data from hundreds of customer Web sites worldwide in such industries as computers and software, education, energy, entertainment, finance, healthcare and pharmaceuticals, insurance, retail and telecommunications. Today this data offers organizations insight into the impact complex Web infrastructures can have on the end-user experience. More important, this data challenges several long-held industry beliefs about Web performance that are considered to be standards by many organizations.

ACTIVEWATCH METHODOLOGY—MEASURING WEB PERFORMANCE FROM THE END-USER PERSPECTIVE

ActiveWatch, based on Mercury Interactive’s Topaz™ application performance management solution, is able to monitor site performance at the application level, from the perspective of an end user. As a result, the data gathered by this service offers insight into how users experience application performance. This is in contrast to the data provided by typical network and systems management tools, which focuses more on the performance of components.

Before monitoring begins, the ActiveWatch team must create scripts that emulate users conducting typical transactions on that site—from logging in as a registered user to the steps taken to transfer money between accounts. These synthetic transactions are then run at predefined intervals, usually every 15 minutes, against the site to verify performance and availability.

Working with the customer, the ActiveWatch team also sets “thresholds” for acceptable Web application performance for different transactions, as shown in Figure 1. When performance falls below these thresholds, customers are automatically alerted via pager, e-mail or an SNMP trap message to an enterprise console. As a result, they have advance warning of a problem and can begin resolution efforts.

Service Level Thresholds

- Service Level Thresholds determines how Topaz reports on the performance of your system.
- Enter thresholds below for transaction performance (response time in seconds).
- The levels that you set will appear on Topaz reports with the following colors:
 OK Warning Poor

TRANSACTION	OK	WARNING	POOR
Home_Page	Less than 4 Seconds	Between 4 and 8 Seconds	Greater than 8 Seconds
Camping_Gear	Less than 4 Seconds	Between 4 and 8 Seconds	Greater than 8 Seconds
Find_a_Store	Less than 4 Seconds	Between 4 and 8 Seconds	Greater than 8 Seconds

Transaction response time is the time it takes for a transaction to be completed.
Response time thresholds are for completed transactions ONLY. Failed transactions are not considered.

Fig. 1. ActiveWatch enables customers to set “OK” and “warning” and “poor” performance levels for each application transaction they want to manage.

ActiveWatch agents located at Mercury Interactive's points of presence worldwide collect performance metrics by running these synthetic transactions on the customer's Web site. Using these metrics, the ActiveWatch team can obtain an accurate reading of the site's performance from the end-user perspective and quickly detect any problems.

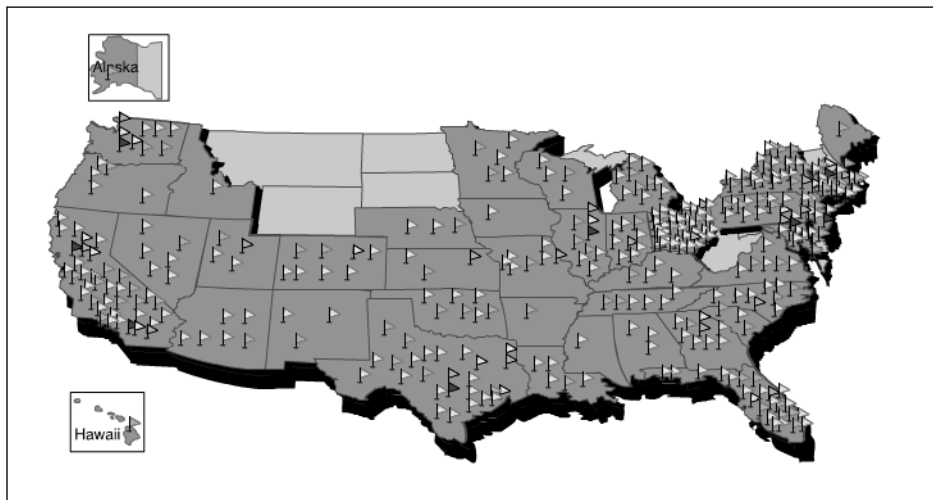


Fig. 2. ActiveWatch can monitor Web application performance from more than 500 locations throughout the world. Within the United States, the most common points of presence for monitoring New York, Newark, Boston, San Francisco and San Jose.

Once any performance problems are identified, the ActiveWatch team can then use these performance metrics to diagnose the root cause. As a result, organizations that use ActiveWatch can resolve Web application availability and performance problems before they become apparent to end users.

REDEFINING WEB APPLICATION PERFORMANCE

LOOKING BEYOND PAGE DOWNLOAD TIME

The application performance experienced by end users is the result of complex interactions among multiple infrastructure components. The more complex the online business process, the greater the number of components within the network, servers and the application that must interact correctly to ensure performance. Users, however, are typically unaware of what needs to happen behind the scenes in order for them to view a home page, log in as a registered user or trade a stock. They are well aware, though, of any performance problems that may result from infrastructure bottlenecks.

For end users, performance is judged in many ways, not just page download times as implied by the 8-second rule. Correct content delivery and availability are also critical. When any one of these aspects of a Web application fails to function properly, performance is considered poor. Too often, standard network and systems management tools fail to detect all of these performance problems because they do not monitor from the end-user perspective. As a result, components might be up and running, but end users are experiencing serious performance problems.

According to Mercury Interactive's ActiveWatch team, end users experience Web application availability problems in three key ways:

- **Application availability.** This can be measured in terms of the user's ability to view the home page, download one or more pages and successfully complete transactions. Application availability is at its lowest and most costly when the site does not work—a transaction, page or the entire site could be down. More complex transactions are prone to lapses in availability, since they rely on more systems interacting successfully. For example, when a user is logging in as a registered user or completing an interactive transaction on the site, many systems must function together correctly to guard against unavailability.
- **Transaction response time.** The ActiveWatch findings show relatively wide ranges in the average times required to load different types of pages on real-world Web sites. This means that some users are having to wait too long for page downloads. These users will often venture to a competitor's site or cause organizations to incur greater operations costs. For example, when users do not use a bank's online application and opt instead to visit the branch, the bank is losing any of the gains in efficiency provided by its Web application investment.
- **Correct content delivery in multi-step transactions.** Even if a site is "up" and response is good, only if the correct content is being delivered will a user have a good experience. If they get an "unable to log on" message or find incorrect information in their banking account, they will likely take their business elsewhere or call their financial institution, which is a less cost-effective way for a bank to handle customer requests.

The challenge facing IT and operations groups today is finding the most effective and efficient means to detect, diagnose and solve Web performance problems in order to provide the best possible user experience. The ActiveWatch data shows that this requires looking beyond industry conventions and rules of thumb for defining acceptable Web performance. Instead, organizations need to focus more on providing a positive end-user experience and setting performance thresholds commensurate with their business objectives.

NEW INSIGHT INTO WEB APPLICATION AVAILABILITY

Many organizations measure availability by server uptime. The industry perception is that if a server is up 99.9% of the time, the Web application must be performing correctly and providing its users with 24x7 availability. The ActiveWatch data shows that this is contrary to reality. Users are experiencing poor performance, even when these organizations claim that they are operating their Web applications at maximum levels of availability.

Before reviewing these findings, it is important to define availability as it is measured by ActiveWatch. Availability is based on three categories that commonly impact the user experience:

- **Page availability.** Delivering the correct page content to the end user within less than 120 seconds. (Most ActiveWatch download attempts were subject to this default timeout period, which can be modified by the customer.)

- **Transaction availability.** Completing the transaction, including returning data appropriate to the transaction, immediately or within the timeout period.
- **User-perceived availability.** Performance based on users' perceptions. For example, performance can be so poor that users will regard a page or a transaction on the Web site as "effectively unavailable."

Based on this definition, all U.S. Web site pages and transactions that were monitored by ActiveWatch exhibited an average availability of 96%—well below the 99+% availability sought by IT managers. Moreover, when these U.S. Web sites were grouped with all sites worldwide, the average worldwide availability was calculated as 92.5%.




Fig. 3. ActiveWatch data revealed average U.S. Web site availability of only 96% and worldwide availability of only 92.5%.

The ActiveWatch data revealed that for Web sites worldwide:

- 20% demonstrated availability greater than 99%
- 57% showed availability of 95% to 99%
- 23% recorded availability below 95%

The 95% figure is equivalent to 1 in 20 transactions failing, which can have a substantial impact on users conducting complex business transactions requiring multiple steps and page downloads.

Performance Thresholds and Availability

Of the millions of transactions they monitored, the ActiveWatch team found that 37% of the time site pages and transactions did not exhibit performance levels considered acceptable. This finding can be broken down as follows:

- Only 63% of ActiveWatch monitoring results showed "OK" performance, defined as download or transaction times better than the customer's pre-set threshold for "warning" performance.
- The other 37% of site pages and transactions exhibited both "poor" Web site performance and "warning" conditions as follows:
 - 15% showed "poor" performance, defined as worse than the customer's pre-set "warning" performance threshold.
 - 22% showed performance "warning" conditions, defined as worse than the pre-set "OK" performance threshold but better than the "poor" threshold.

One function of the “warning” threshold setting is to alert IT teams to degradations in performance that need to be addressed immediately, before they impact users. The second function is to help make IT aware of other performance degradations that are not critical or noticeable to the user at that moment but that may escalate over time. Using this information, IT groups can fine-tune servers or network devices within the infrastructure to resolve these problems and bolster performance.

According to ActiveWatch data, up to 15% of all transactions monitored generated such alerts, or 1 out of every 7 transactions. In addition, for 72% of the Web sites monitored, at least 10% of the performance measurements fell short of the firms’ pre-set performance thresholds. Moreover, for almost 75% of all sites, at least 10% of users (1 out of 10) experienced performance that was worse than the threshold that the ActiveWatch customer had pre-set as unacceptable.

DETERMINING ACCEPTABLE PERFORMANCE LEVELS

Page Download Time v. Transaction Time

Over the past year or two, the 8-second-rule has become the yardstick for measuring acceptable performance. It states that Web site users will click away from a Web site page after waiting 8 seconds for a page download. ActiveWatch performance data challenges this notion. More important, the ActiveWatch data shows that organizations cannot think about Web site performance solely in terms of page download times, but must look at transaction times. As a result, they need to set realistic thresholds for their users.

By analyzing the ActiveWatch data, Mercury Interactive’s experts uncovered three common transactions on most Web sites: download time for the home page, the time for an end user to log in to their account as a registered user and search functions. After monitoring a wide range of customer sites, ActiveWatch experts found:

- Average home page download time: 4 seconds
- Average login time: 12 seconds
- Average search function time: 10 seconds

The ActiveWatch data shows that there is room for variance in performance thresholds, and that a rule-of-thumb or 8-second rule approach is not the most effective. Rather, organizations need to set performance thresholds based on their users’ business needs and expectations. If users feel the data is critical, they may not mind waiting longer than 8 seconds for a transaction provided it is completed correctly.

Thresholds then can vary substantially by function. For example, while users expect to see a home page—complete with a reasonable number of graphical elements—in just 4 seconds, they will wait longer than 8 seconds for other key functions.

Other Factors Affecting the End-user Experience

The ActiveWatch team was also able to isolate other factors that can hinder Web performance, as well as suggest remedies for them.

Static v. dynamic pages. In addition to home pages, many customers directed ActiveWatch to monitor the performance of pages inside the site, including static and dynamic pages. ActiveWatch experts analyzed data collected for text-based static pages—pages without graphical elements. According to ActiveWatch data, these static pages loaded, on average, in 1.5 seconds, compared to 3.5 seconds for dynamic pages.

A Web site team then can conceivably double the performance of Web pages by converting them from dynamic to static pages. This course of action is appropriate when speed is more important than serving dynamic content. However, if the dynamic content is adding true value to the user, he or she will accept the longer download time. This example further illustrates why performance thresholds should be guided by function and user expectations.

Secure v. non-secure pages. ActiveWatch data also confirmed that pages transmitted over secure connections generally performed more slowly than non-secured pages. ActiveWatch's root-cause analysis data further revealed that secured pages often are implemented before a user logs in as a registered user. Since there is no need for security functions before a user logs in, one way to easily improve performance is to change these secured pages into non-secured pages.

ROOT-CAUSE ANALYSIS: DIAGNOSING PROBLEMS

Identifying Web application availability issues is only the first step in the process of Web performance management. The second, and often far more difficult, step is diagnosing the root cause within complex Web infrastructures. ActiveWatch's sophisticated root-cause analysis tools enable organizations to quickly and accurately pinpoint the source of performance problems.

INTERMITTENT PROBLEMS

Because they appear irregularly, intermittent problems can be the most challenging to diagnose. Often in such cases, a user calls to complain of a problem and the organization's help desk is both unaware of the problem and unable to reproduce it. In fact, organizations seldom can detect or diagnose the cause of intermittent availability problems that extend beyond server downtime.

ActiveWatch, however, enables organizations to systematically record all intermittent performance problems from the end-user perspective and diagnose their root cause within the infrastructure. This is particularly critical, since the ActiveWatch data revealed that 70% of all unavailability and performance conditions can be classified as intermittent.

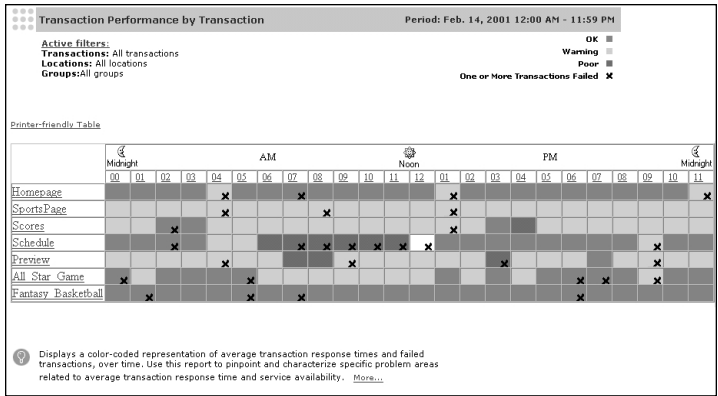


Fig. 4. Data displays for customers of the ActiveWatch service reveal intermittent performance problems.

USING ROOT-CAUSE ANALYSIS TO DIFFERENTIATE BETWEEN SERVER AND NETWORK PROBLEMS

Not only does ActiveWatch enable users to determine the source of performance problems, but it has provided enough data to challenge the industry belief that most performance problems originate within the network. The data shows that 75% of problems with Web application availability and performance actually are server problems, and 25% originate within the network. Server problems found by ActiveWatch include bad session IDs and inefficient caching. Typical network problems include poor peering arrangements and router configuration problems. ActiveWatch’s root-cause analysis tools enable IT groups to pinpoint the precise source of problems within the network, the server or the application itself for more efficient problem resolution.

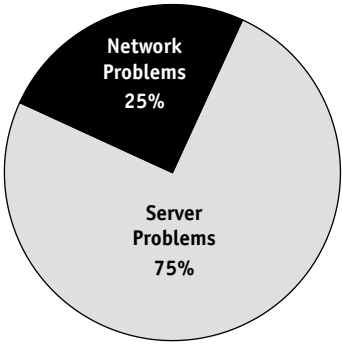


Fig. 5. ActiveWatch revealed that far fewer performance problems were in the network than in the server.

SUMMARY

Mercury Interactive's ActiveWatch team has collected valuable data on the ways complex Web infrastructures impact the end-user experience. This data, collected over a period of 18 months, challenges many of the industry's long-standing beliefs about Web performance. More important, it enables organizations to develop more effective Web application performance management programs and thereby deliver better performance to their end users.

ABOUT MERCURY INTERACTIVE

Mercury Interactive is the leading provider of enterprise testing and performance management solutions. The company's automated software and managed services help companies deliver and maintain high-performance applications. Customers worldwide use Mercury Interactive solutions across their application and technology infrastructures to minimize hardware and operational expenses, protect revenue streams and enhance their competitive positions.

Mercury Interactive was founded in 1989 and is headquartered in Sunnyvale, California. The company has over 1400 employees with offices in more than 20 countries. For more information on Mercury Interactive, visit the company's Web site at www.mercuryinteractive.com.