

WebSpy: An Architecture for Monitoring Web Server Availability In a Multi-Platform Environment

(Technical Report CS-2002-07)

Madhan M. Thirukonda
&
Shirley A. Becker
Computer Science Department
Florida Institute of Technology
150 West University Blvd.
Melbourne, Florida 32901
{mthiruko, becker}@cs.fit.edu

Keywords: Web Server, Web Server Management, Web Server Monitoring, Web Server Availability, Web Server Downtime, Electronic business, Electronic commerce, Enterprise Web Server.

Abstract

For an electronic business (e-business), customer satisfaction can be the difference between long-term success and short-term failure. Customer satisfaction is highly impacted by web server availability, as customers expect a web site to be available twenty-four hours a day and seven days a week. Unfortunately, unscheduled web server downtime is often beyond the control of the organization. What is needed is an effective means of identifying and recovering from web server downtime in order to minimize the negative impact on the customer. An automated architecture, called WebSpy, has been developed to notify administration and to take immediate action when web server downtime is detected. This paper describes the WebSpy architecture and differentiates it from other popular web monitoring tools. The results of a case study are presented as a means of demonstrating WebSpy's effectiveness in monitoring web server availability.

Copyright 2002 by the authors. All rights reserved. Please do not print, copy, or distribute without express permission of the authors.

1. The Need for an Effective Web Server Monitoring Tool

There is no question that the availability of a web server plays a major role in the success of e-business. Ideally, a web server should function as expected anytime, anywhere, and for any customer. The reality is that unscheduled web server downtime happens, and often times it is due to factors beyond the organization's control (Schmidt, 2000). A disgruntled customer has a ripple effect on the use of a web site in that many other customers are told about the negative experience. These customers, in turn, disseminate information to other customers causing long-term and perhaps irreparable damage to the organization. Proactive measures need to be taken to maintain a consistent, high-rate of web server availability. An effective means of monitoring a web server is needed such that downtime can be minimized.

An innovative web monitoring architecture, called *WebSpy*, has been developed with this in mind. The automated tool, which supports this architecture, has special features including dynamic monitoring of a web server, disseminating timely information, and maintaining the status on availability. It was built to satisfy the web server monitoring requirements of a large e-business organization, which couldn't be met by existing commercial tools.

The enterprise web server of the organization, for which *WebSpy* was built, deploys server-side Java technology using third party application servers to serve its web users. Many enterprise web servers are increasingly using application servers in order to take advantage of dynamic applications that run on the server side of an e-business. An application server interacts with the database server to get the business data required by the end users. Due to the complexity of such an architecture, it is very important to ensure high availability of the web server.

It is proposed in this paper that WebSpy be used to effectively minimize the unscheduled downtime of an enterprise web server. Section 2 presents an overview of World Wide Web (web) concepts. Section 3 describes the architecture of WebSpy tool from the point of view of an administrator. Section 4 describes the functionality of the tool and the protocol used to implement the architecture. Section 5 summarizes a case study used to evaluate the capabilities of the tool in meeting the requirements specified by the e-business. The paper concludes with future research directions.

2. World Wide Web (Web) Concepts

The Internet is a network infrastructure built on certain standards that provide the means for users to communicate with one another. The web is one of the services offered by the Internet, and its architectural framework allows users to access geographically dispersed, linked documents (Tanenbaum, 1996). The web also provides the capability for a running a multitude of applications using three simple mechanisms (Ciancarini, et al., 1996):

1. URLs (Universal Resource Locators) provide information on the location of a document and the protocol needed to access it.
2. A client-server protocol, called Hyper Text Transfer Protocol (HTTP), provides the means to fetch a document from a web server.
3. A simple markup language, Hyper Text Markup Language (HTML) allows content authors to describe the information and store them in a document. A HTML document is also referred to as a web page.

A main component of the web is a web browser, which is a program that interprets and displays a web page using HTML specifications. The most popular web browsers are Netscape

Communicator and Internet Explorer though there are others. The other main component is the web server, which is a program that delivers web pages upon a HTTP request from a web browser. Fielding *et al.* (1999) define the HTTP protocol as a request/response protocol. A web browser requests a web page by sending the HTTP request message that contains the URL of the web page to the web server. The web server responds with a HTTP response message followed by the web page. E-business enterprises often use application servers to complement the web server.

Table 1: Factors Affecting Web Server Availability	
Factor:	Explanation:
Server software failure	The inability of the web server software to serve a request possibly due to its bad run time environment (e.g., memory leak).
Support server software failure	Supporting software (e.g., application or database server software) fails such that the web server cannot respond to a HTTP request.
External factors	External factors (e.g., hacking attempt) impact the availability of the web server.

In general, availability is defined as the percentage of time during which a system is available for a given purpose (CAIDA www.caida.org, 2001). We define web server availability as the percentage of a time interval during which the web server is available to serve HTTP requests. Web servers may fail to serve HTTP requests due to a list of possible factors, as presented in Table 1. On an immediate basis, the e-business is focused on getting a web server that is unavailable up and running as quickly as possible. In the long run, it has a goal of maintaining a consistently high rate of web server availability.

Ideally, an e-business would want to establish an availability rate for each of its web servers, and then gather statistical information about whether this rate is achieved. It would also want to track fluctuations in the availability rate so it can investigate the causes of availability problems. Thus, meaningful information needs to be captured at the time of failure in terms of

down time, possible cause of failure, and recovery mechanisms. To support these goals, a process is needed to periodically check the availability of the web server and take necessary action in the event of a problem. The process would alert a system administrator as soon as a problem is detected. In some cases, a corrective action is automatically taken as soon as a problem is detected, thus preventing human intervention. Such automatic recovery guarantees quick restoration of web server availability.

3. The WebSpy Architecture

Table 2 identifies the basic features of commercial web server monitoring tools. Though the commercial tools¹ that were assessed as part of this study met these basic requirements, none met the additional monitoring needs of the e-business as specified in Table 3. These additional monitoring requirements, typical of many of today’s e-businesses, were the motivation for the development of WebSpy.

Table 2: Basic Features of Commercial Web Server Monitoring Tools	
Feature:	Description:
Periodic monitoring	The web server is monitored on a schedule based on its availability requirements. Basic timer settings include “sleeping time,” the interval between two monitoring cycles, and “slow response timeout interval,” the threshold time within which a response should be obtained from the web server.
Alerts	An alert mechanism notifies individuals when the web server is unavailable. An email note, for example, is sent to a designated administrator.
Log entry	A log entry is made when the web server is unavailable in order to support the calculation of availability statistics.

¹ Web monitoring solutions, used in this study, are the same ones cited in Stanford Linear Accelerator Center of Stanford University (Cottrell, 2001).

Table 3. Additional Requirements of the E-business		
Requirement:	Description:	Commercial Support:
Multiple Operating System Platforms	The monitoring service runs on multiple platforms to support the production web servers that run on multiple platforms.	None
Start as NT Service	The monitoring tool is automatically started as a service in the Windows NT environment	Limited
Additional Timer Settings	The monitoring service starts monitoring only after the web server has been started. An initial timing delay is needed to allow for the web server to start up. This “initial sleep time” eliminates false alarms associated with start up when the web server host is rebooted. The time between retry attempts (referred to as “retry time”) is specified by the administrator.	Limited
Web Interface for Viewing Log Data	A web interface provides viewing capability of the log information for all web servers.	Very Limited
Recovery Scripts	The monitoring tool is capable of executing a recovery script when the web server becomes unavailable. The script has commands to restart the web server, the database server, or the web server host depending on the problem. The administrator configures certain commands in the recovery script that execute when a specific error occurs (the error message is displayed as part of the web page log information).	Very Limited
Secondary Email Alerts	If the web server is unavailable for a specified time, then a secondary email list is used for notification purposes. This provides the capability of alerting officials and managers of prolonged problems in web server availability.	None
Availability Metrics	Web server recovery timestamps provide useful information in calculating availability metrics.	None

Figure 1 shows the WebSpy architecture comprised of two parts; the WebSpy station where both the web server and WebSpy are installed, and the workstation where administrators receive alerts and view availability reports. WebSpy has two major components, the WebSpy client (hereafter referred to as client) and the WebSpy server (hereafter referred to as server). Each of these components is explained.

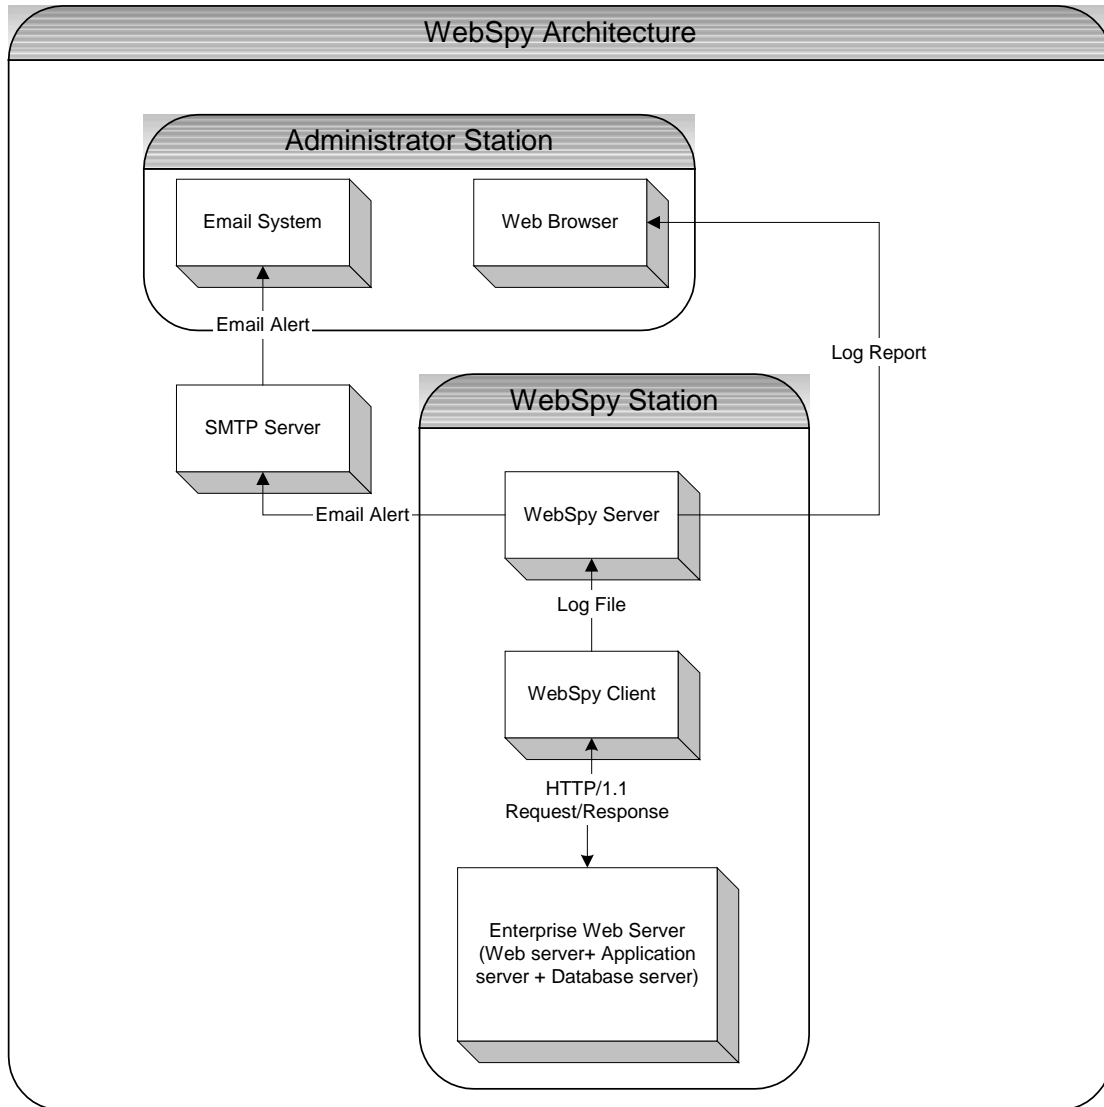


Figure 1. Enterprise Web Server with Monitoring Ability

3.1 Client

The client is a program that periodically monitors the web server availability based on settings specified by the administrator. The client is installed on the same computer as the web server. Table 4 lists the settings that are the basis for monitoring. WebSpy allows these properties to be dynamically updated for effective monitoring. The client can concurrently monitor several web servers installed on the same web server host. For each web server, a list of settings is specified.

Table 4: Settings for Web Monitoring	
Properties:	Description:
URL	URL to be monitored (e.g., <i>http://cs.fit.edu/profiles.html</i> has components; <i>HTTP</i> protocol, host name <i>cs.fit.edu</i> , path for the resource <i>profiles.html</i>). One URL should be specified for each web server installed on a web server host.
Enable	The monitoring of a web server is enabled using this property.
Init_time	An initial time is used to put the client to sleep in order to synchronize the startup of the web server and the client.
Sleep_time	A sleep time is used as a time interval in which the web server is not monitored. It needs to be reasonably short in order to detect availability problems in a timely manner. It also has to be reasonably long to minimize the impact on resources and overall server response time.
Slow_time	A slow time is used to specify a maximum time period for receiving a server response.
Retries	The client checks for web server availability a specified number of times, after which a failure is declared.
Retry_Interval	The time in seconds between two clients retry attempts in determining server availability.
Text_failed	The administrator can customize text that becomes part of the email alert.
Script_on_failure	A recovery script that is automatically executed when a server is unavailable. The script may have several commands based on the command line arguments passed to it. WebSpy executes the script by passing the command line arguments depending on specific error message that appear on the web page.
Email	Alert messages are sent to the comma-separated list of email addresses.
Log_transfer_interval	The log transfer interval specifies when the client sends the log file to the server.
Dynamic_Errors	A list of error messages of which one or more may appear in the web page if there is a problem in dynamic web page generation. Based on a particular error message, an appropriate command in the recovery script would be executed.
Escalation_Time	The time in minutes after which any problem in availability should be escalated.
Secondary_Email	Alert messages are sent to this comma-separated list of email addresses, after Escalation_Time. Email addresses of managers are typical for this category.

The client assesses the availability of a web server by monitoring whether the web server is successful in serving a HTTP request. The HTTP request is based on a URL specified by the

administrator. The web page specified by the URL may be a static or dynamic one.² It is important that the URL be valid and virtually free from programming errors.

A log file is used to store information generated by the client while monitoring the server. This information includes availability status, email alerts, and recovery status, among others, as shown in Table 5. This information is useful for calculating availability statistics for each web server that is being monitored.

Table 5: Collective log file information	
Parameter:	Description:
EventID	Each WebSpy event has a unique number associated with it. (WebSpy events are discussed in detail in Section 3.3)
URL	URL being monitored
Response	The HTTP response header information provided by the web server
Fail Time	Local time when a problem in availability was detected
Recovery Status	Status of the recovery as executed by the client
Email Identifier	List of registered email addresses for receiving email alerts
Mail Status	Status of an email alert
Mail Content	Content of an email alert
Alert Message	Alert message formatted to be displayed in a web browser
Host	IP address of the client
Server Name	Name of web server software (includes version information)
Recovery Time	Local time of recovery

In addition to monitoring a web server, the client is responsible for executing a script for restoring the availability of a web service. The recovery script is executed with a command line argument based on the error message on the web page. The recovery script should be configured to restart appropriate servers based on the command line argument. The output of the script is updated to the log file.

² In e-business production web servers, web pages are generally expected to be dynamic (created when the request is made) since the data is highly volatile and constantly changing. Often, the creation of dynamic web pages requires the service of an application server and a database server.

3.2 Server

The architecture of WebSpy supports multiple instances of the client in order to monitor different web server hosts. Only one server is needed to manage these client instances, and it can reside on any of the web server hosts or on a separate workstation.

The clients periodically communicate with the server to send it the log information. The server collects the log information of each client instance and stores it in a cumulative log file. The server uses an internal Java servlet engine to periodically generate a web report from the monitoring results stored in the cumulative log file. This servlet engine runs on a different TCP/IP port than any other TCP/IP ports used by the enterprise web server. The report, shown in Figure 2, is listed as a URL in the email alert sent to the administrator. The server requires access to a Simple Mail Transfer Protocol (SMTP) server to forward email alerts.

WebSpy Report and Alert			
URL	Failure Information	Time Failed	Recovery Information
http://163.118.134.52/hello.jsp	Web service unavailable (No Response).	Thu Jul 26 11:35:44 EDT 2001	Site recovered from failure (on Thu Jul 26 11:36:47 EDT 2001).
http://163.118.134.52/hello.jsp	Application server error (Code=503, Message=Could not connect to JRun).	Thu Jul 26 11:43:56 EDT 2001	Site recovered from failure (on Thu Jul 26 11:44:58 EDT 2001).

Alert Message from WebSpy (163.118.134.68) => FIT CS Website is inaccessible
URL "http://cs.fit.edu" (URLName: "FIT CS Website") failed on Sat Apr 27
17:07:07 EDT 2002 after 0 retries.
Monitor Status => Web service unavailable. (No Response).
For WebSpy web interface to log information, please visit
<http://163.118.134.68:8002/servlet/Reporter>

Figure 2. WebSpy Report and Alert

3.3 Events

During the monitoring process, one of four events could occur that require an action taken by the client. For each event that occurs, the information describing it is written to the log file. Table 6 lists each event and the information stored for it.

Table 6. WebSpy Events	
1) Failure After Retry Attempts - Event fired after the number of retry attempts is equal to the “retries” parameter value and the web server is still unavailable.	
Attribute	Value
EventID	1
Response	The HTTP response header information provided by the web server, if received.
Server_Name	Web server name, if received.
2) Email Notification – Subsequent event that is fired if any email address is registered to receive email alerts.	
Attribute	Value
EventID	2
Email_ID	Email address registered with the client.
Mail_Status	Null (The status of email alert is stored after its delivery).
Mail_Content	Email alert message.
Alert_Message	Email alert message formatted to be displayed in a web browser.
3) Execution of Recovery Command – Event that is fired after executing the recovery script.	
Attribute	Value
EventID	3
Recovery_Status	The output and the return code of the command execution.
4) Web Server Recovery – Event that is fired when the web server is recovered after a failure has occurred.	
Attribute	Value
EventID	4
Recovery_Time	Local time of the client when the web server is recovered.
5) Email Notification: Escalated – Event that is fired when the web server is still unavailable after Escalation_Time.	
Attribute	Value
EventID	5
Email_ID	Email address registered with the client.
Mail_Status	Null (The status of email alert is stored after its delivery).
Mail_Content	Email alert message.
Alert_Message	Email alert message formatted to be displayed in a web browser.

4. WebSpy Functionality

The main functionality of the client is to monitor the web server and take appropriate actions when the web server is not available. The main functionality of the server is to act as a coordinator in collecting the monitoring reports from each client, sending email alerts, and making the cumulative report available online.

4.1 Client Functionality

Once activated by an administrator, the client waits the specified *Init_Time* time period, and then sends the web server a HTTP request using the specified URL. The client waits for a HTTP response from the web server. If a response header is not obtained within the specified *Sleep_Time* time period, then the client repeatedly sends the HTTP request until the number of retries is equal to the *Retries* parameter value. Each retry attempt occurs after the specified *Retry_Interval* time period has lapsed. If no server response is obtained during or after the specified *Retries* attempts, the web server is considered not available and the log file is updated.

When a HTTP response is received, the client extracts the response code and message from the header. The *2xx* response code represents a successful transaction, which is an indication that the web server is available. The web server is considered unavailable if there is no response or if a *5xx* response code is received. *5xx* response code means that the web server was unable to serve the HTTP request because of a server-side problem.

To illustrate this, we use an example from our web monitoring experiences. A recurring problem encountered during the analysis stage of the selected architecture was denoted by a *503* error code. This is defined in the HTTP specification as a *Service Unavailable* problem (Fielding *et al.* 1999). This problem may occur when the web server is unable to receive services of another external application. For our web architecture, JRun 3.1 was the application server. The client received an error code *503 JRun closed connection* when the JRun application server was not purposely stopped.

The final determination of web server unavailability occurs only after executing the specified retry attempts. In the case of a 5xx response from the web server, the response message is compared with the list of error messages called *Dynamic_Errors*. In this list, the administrator would include messages that he suspects appear on the response message due to a problem during dynamic page generation (in our example, an error message in the e-business web server was 'JRun closed connection').

```
Recovery Script

@echo off
REM Script to run upon web server failure

REM Based on the command line parameter %1, execute different commands
if %1 == 1 goto appl_server
if %1 == 2 goto web_server
if %1 == 3 goto server_host
goto exit

:appl_server
REM Command to restart application server
goto exit

:web_server
REM Command to restart web server
goto exit

:server_host
REM Command to restart server host
:exit
```

Figure 3. Recovery Script

When a match is found in the list, the failure is reported as an *Application server error* with the response code and message, and then the client executes the recovery script. The command line argument for the recovery is calculated based on the position of the error message in the list. The recovery script is written to execute the appropriate commands based on the command line argument as shown in Figure 3.

When there is no matching code or no response from the server, it is reported in the log as *Web service unavailable*. WebSpy events, presented in Table 5, are fired and the log file is updated. The client continuously monitors the URL until the web server becomes available. The log files are frequently sent to the server based on the parameter *Log_transfer_interval* (as presented in Table 4).

4.2 Server Functionality

The server listens to a specified TCP/IP port for client connection. It receives periodic connections from the clients and gets the log information. On receiving the log file from the client, the server looks for the *Email notification event* record (event 2 in Table 6). If it is found, then an email is sent based on the information and the email alert status is stored in the cumulative log file. Other event information is also stored in the cumulative log file, which is later used to produce the cumulative report. The server has a Java servlet engine to dynamically generate a report based on the information in the log file. The Java servlet that is used to generate the report can be invoked by using the URL, which is contained in the alert sent to the administrator.

4.3 Performance Impact

The HTTP transactions, due to WebSpy monitoring, have an impact on web server performance because they simulate user requests. Reducing the number of transactions assessing availability minimizes the negative impact on performance. The parameter *Sleep_Time* defines the time during which the client would wait to send a HTTP request. *Sleep_Time* should be reasonably set to detect availability problems as quickly as possible without causing a noticeable performance impact on the web server.

A worst-case analysis is described in terms of the performance impact of WebSpy monitoring the web server. If the administrator enters sixty seconds (minimum allowed limit) as the value for the parameter *Sleep_Time*, then the total time required for a monitoring cycle is calculated as follows:

$$T_{mt} = T_{et} + T_{st}$$

Where: T_{mt} is the total time for the monitoring cycle, T_{et} is the execution time of a HTTP transaction for a URL request, and T_{st} is the sleep time required for the monitoring cycle.

For illustration purposes, let's assume that T_{et} takes a minimal amount of time. Thus, $T_{mt} \sim T_{st}$. This means that there is approximately one HTTP request per minute. For a twenty-four hour time period, there are approximately $24 * 60 = 1440$ HTTP requests from the WebSpy monitoring tool.

5. Case Study

A case study was conducted to determine the effectiveness of WebSpy in monitoring the availability of a web server given the e-business requirements. The web server architecture was comprised of Internet Information Server (IIS) as the web server and JRun as the application server.

The initial sleeping time *Init_Time* was set to 20 seconds, the interval between two monitoring cycles *Sleep_Time* was set to 60 seconds, and slow response time out interval *Slow_Time* was set to a maximum of 45 seconds. These parameter settings are typical of the e-business' real-world setting for minimizing false alarms (web server is available but reported have failed) and maximizing availability by reporting problems quickly after they occur (Loren Buhle *et al.* 1996). The client was set to transfer the log file to the server periodically with an

interval of 30 seconds. A recovery script was written such that it would restart the web server for a '1' command line argument, restart the application server for a '2' command line argument, and restart the web server host for a '3' command line argument.

A program was written to simulate users requesting HTTP service of the web server. The web server had twenty URLs, ten of which were static HTML pages ranging from 17 Kbytes to 234 Kbytes in size. The other ten URLs were dynamically produced resulting in a HTML page ranging from 1Kbyte to 350 Kbytes in size. Simulated (virtual) users were set to randomly download any one of these 20 web pages (every random time between 1 and 5 minutes to simulate reading).

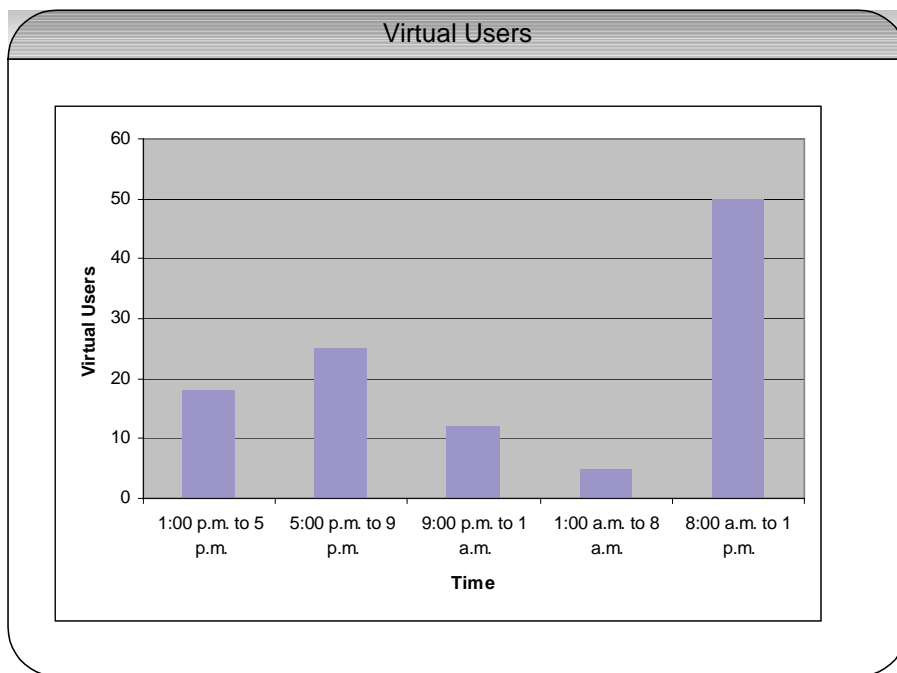


Figure 4. Virtual Users Traffic

The maximum number of user requests (clients) for the web server was restricted to ten so that the maximum limit was reached thus ensuring the web server was busy. Figure 4 shows the number of users per time interval in the simulated, twenty-four hour time period. The number of users peaked at fifty during the morning hours, as is generally the case for many e-

businesses. To simulate the web server's inability to serve a HTTP request, we stopped the application server and the web server once for a several minutes.

5.1 WebSpy Behavior

WebSpy detected the availability problem that occurred during the case study and restored the web service as expected. As soon as a problem in web server availability was detected, WebSpy sent an email to the specified administrators. The *Web Interface* component of the WebSpy displayed detailed information including failure detection time, web server name, status of the restoration command executed, recovery time, and other pertinent information. For each failure, WebSpy detected it and restored web service immediately as evidenced in Figure 5. The *Web Interface* provides the information obtained from the log file. Each hyperlink in the *Web Interface* is linked to another web page that provides detailed information on the output of the recovery script executed, status of the email alert, and a copy of the email alert message.

WebSpy Report			
URL	Failure Information	Time Failed	Recovery Information
http://163.118.134.52/hello.jsp	Web service unavailable (No Response).	Thu Jul 26 11:35:44 EDT 2001	Site recovered from failure (on Thu Jul 26 11:36:47 EDT 2001).
http://163.118.134.52/hello.jsp	Application server error (Code=503, Message=Could not connect to JRun).	Thu Jul 26 11:43:56 EDT 2001	Site recovered from failure (on Thu Jul 26 11:44:58 EDT 2001).

Figure 5. WebSpy Report

5.3 Availability Metrics

There are two metrics that provide useful information in terms of web server availability. The *Mean Time to Repair* (MTTR) provides the average time taken to restore normal service after a loss of availability. This metric provides ongoing feedback as to the downtime associated with a server, and its potential impact on customer satisfaction and retention. The other important

metric is *Mean Time Between Failures* (MTBF), which provides the average time between the beginning of normal service and the next loss of availability. This metric also provides insight into the long-term impact of server failures on an organization's customer base. The availability metrics are useful from the administrator point of view in terms of finding the stability of the web server and the application server. Figures 6 and 7 show the *Time To Repair* statistics and the *Time Between Failures* statistics for the web server.

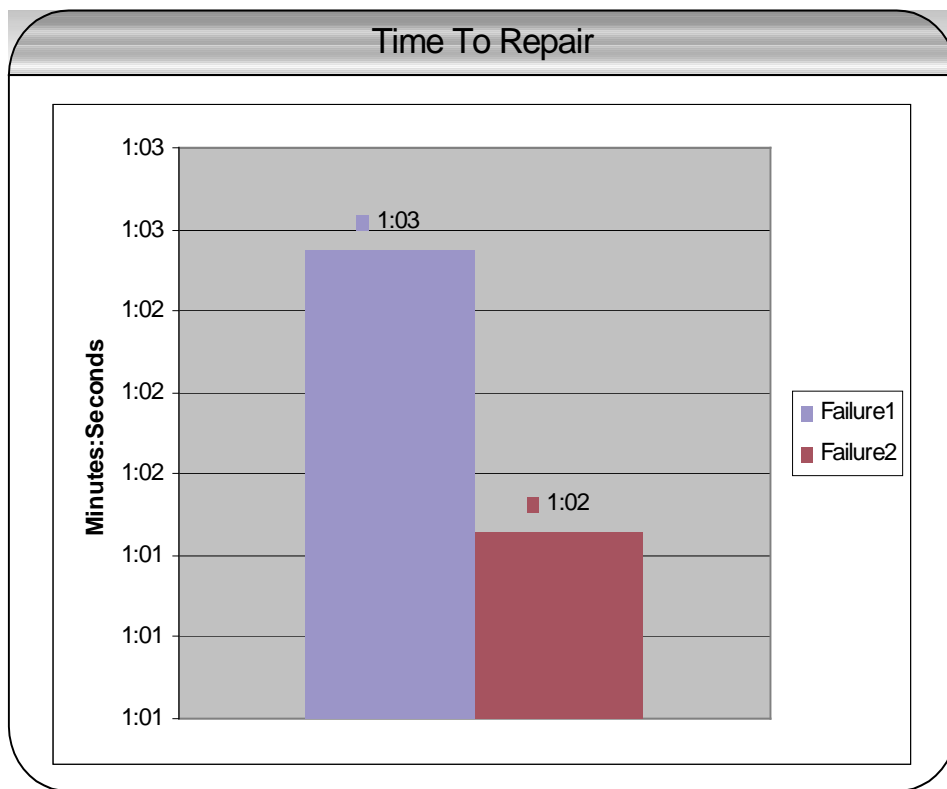


Figure 6. Time To Repair

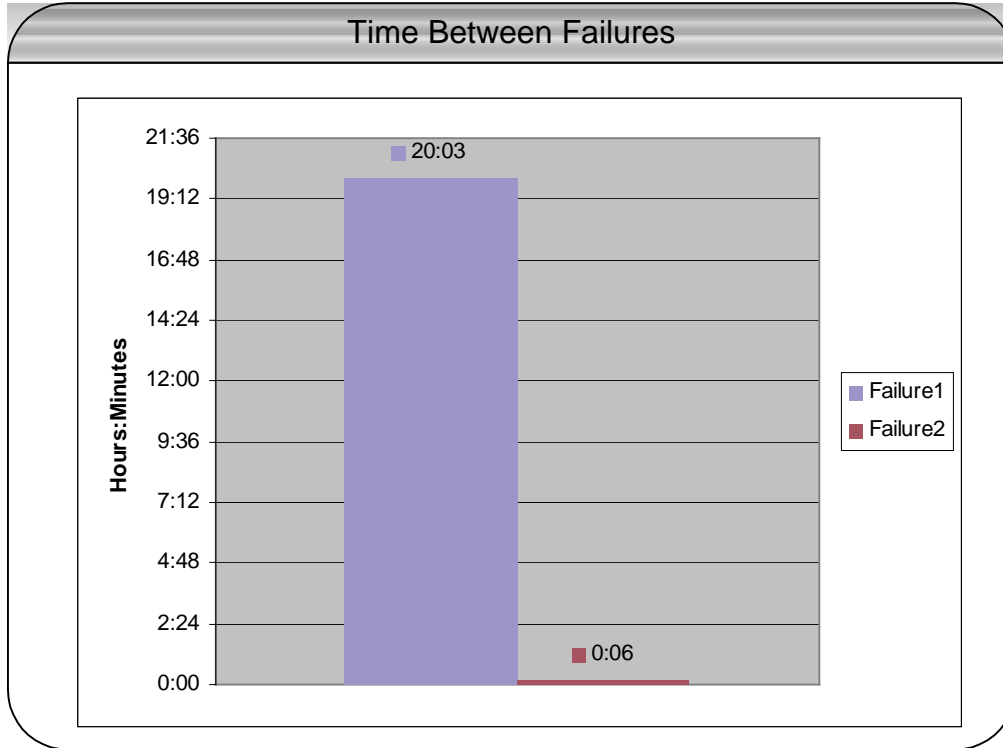


Figure 7. Time Between Failures

The *MTTR* (~1 second) maintained a low value meaning that the web service was restored very quickly. The *MTBF* (~10 hours) was quite high, which means that the web server was stable and not failing at a high rate. For the web server architecture, used in this case study, we can define availability in terms of these two metrics (Blanchard, 1997). We start defining uptime (*MTBF*) and downtime (*MTTR*):

$$\text{Uptime} = \Sigma \text{ Time Between Failures}$$

$$\text{Downtime} = \Sigma \text{ Time To Repair}$$

The uptime for the web server was 86,275 seconds and downtime was 125 seconds. Availability is calculated as:

$$\text{Availability} = \text{Uptime} / (\text{Uptime} + \text{Downtime}).$$

$$\text{Where: Availability} = 86275 / 86400 = 99.85\%$$

What we can conclude from this data is that the web server was maintaining a high rate of availability during the case study, as recovery scripts were executed quickly after a problem detected. The web server was automatically recovered quickly, which means that unscheduled downtime was kept to a minimum.

9. Overall Findings and Future Research

The WebSpy provided several features, as described in Table 7, that are not fully supported by other commercial tools. The email alert capability is extremely useful, as it provides additional information in a web-based environment for tracking web server failures. It also runs in multi-platform environments, unlike any of the other commercial tools. The WebSpy has an automatic recovery capability, which is necessary for fast recovery from certain types of failures.

Table 7. WebSpy Features		
Features	Description	Commercial Tool
Multiple Platform Support	Webspy is built in Java, which virtually runs on all operating systems.	None
Web Interface to Log Information	A single report presents information about the availability of the web servers, which is viewed via a browser.	Very Limited Support
Customizable Automatic Recovery	Recovery scripts are automatically executed thus minimizing human intervention and time delays. Recovery scripts can be customized for different possible problems regarding web server availability.	Very Limited Support
Secondary Email Alerts	Email alerts are sent to designated personnel when there is prolonged unavailability. This is in addition to the primary email list that is notified as soon as a problem is detected.	None
Availability Metrics	Web server recovery timestamps provide useful information in calculating availability metrics.	None

One of the weaknesses of WebSpy is its report generation capability in terms of performance issues associated with web site access. It would be very useful, for example, to maintain statistics on web page load times, average number of retries attempted for web page access, and the average throughput given environmental constraints. This information would prove useful in not only monitoring but also in predicting availability problems. In this way, downtime would be minimized as certain types of problems could be prevented.

There is no doubt that other improvements can be made to the WebSpy and other commercial tools in the marketplace. But, WebSpy offers features that none of the other tools completely support, which makes it very useful for e-businesses working in the global, online marketplace. On a final note, the e-business that funded this research has been using WebSpy successfully for the past six months to monitor various web servers in a geographically dispersed environment. This is a reflection of the capability provided by WebSpy in meeting the needs of a global e-business enterprise.

References

- R. Fielding, UC Irvine, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, *Hypertext Transfer Protocol -- HTTP/1.1*, The Internet Engineering Task Force, Request For Comments (RFC) 2616, June 1999. Electronic version
<http://www.ietf.org/rfc/rfc2616.txt>
- P. Ciancarini, R. Tolksdorf, and F. Vitali, *Weaving the Web in a PageSpace using Coordination*, Technical Report, January 1996. Electronic version
<ftp://ftp.cs.unibo.it/pub/cianca/pagespace.ps.gz>
- CAIDA Metrics Working Group, *Network Measurement FAQ*, Cooperative Association for Internet Data Analysis, Super computer center, University of California, San Diego, January 2001. <http://www.caida.org/outreach/metricswg>
- Benjamin S. Blanchard, *System Engineering Management*, Wiley-Interscience Publication - Second Edition, pp.120-128, November 1997.
- Andrew S. Tanenbaum, *Computer Networks*, Prentice Hall - Third Edition, 1996.
- Loren Buhle *et al*, *Webmaster's Professional Reference*, New Riders Publishing, pp. 468 – 497, 1996.

Les Cottrell, *Network Monitoring Tools*, Stanford Linear Accelerator Center of Stanford University, March 2001. <http://www.slac.stanford.edu/xorg/nmtf/nmtf-tools.html>

Jürgen Schmidt, *To Be Up or Not To Be Up - Analysis of Web Server Downtimes*, August 2000. <http://www.heise.de/ct/english/00/08/174/>