



Proactive Real-time Monitoring of Real Web Transactions

W h i t e P a p e r



Restricted Rights Legend

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations Inc. eG Innovations, Inc. makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

Copyright

© Copyright 2003 eG Innovations. All rights reserved. eGurkha and eG ASPLite are trademarks of eG Innovations. All other trademarks, marked and not marked, are the property of their respective manufacturers. Specifications subject to change without notice



Introduction

As the web has evolved from being primarily an information exchange medium to a place where real business is transacted, it has become critical for IT Businesses to ensure that their web sites remain available twenty four hours a day, seven days a week. Moreover, faced with increased competition IT Businesses have to ensure that their web sites offer the best experience for their customers. One of the critical factors governing the user experience is the responsiveness of any website. Since the user experience is governed predominantly by how well the application components associated with an IT Business site perform, effective and proactive application monitoring that drills down to the root-cause of a problem if and when it occurs, is imperative.

Current Monitoring Approaches

IT Businesses normally use any one or more of the following categories of monitoring tools:

Request Emulators: Request Emulators generate synthetic requests periodically from one or more locations to different web sites. The monitoring can be done from within a data center that hosts the IT Business or from various locations in the Internet (i.e., external to the data center).

The main drawbacks of this approach are:

- All application transactions cannot be monitored by simple emulation of the transactions. For example, it may not be possible to emulate a user making a payment, or a user registering to the site.
- Emulation-based techniques mainly sample the target environment. Hence, if a specific transaction is failing, say 10% of the time, the chance that emulated monitoring is able to detect the problem is only 10%.

Hence, emulation techniques typically detect and report problems only when they are severe enough to impact the end-user performance.

Log File Analyzers: These analyzers periodically process the numerous log file entries produced by the web server, and report a number of statistics including user profiles, access patterns, response times, etc.

The main drawbacks of this approach are:

- A log analyzer generates hundreds of graphs that are difficult to analyze
- For most heavily loaded web servers, the log file analysis process can be time consuming and resource intensive. Hence, log file analysis is mainly used in off-line mode. Consequently, this approach cannot be used for monitoring, detecting, and fixing problems in real-time.

eG overcomes the above drawbacks using its unique web adapter capability, which performs internal monitoring of real web transactions (not emulated ones). The monitoring is done in an implementation-independent manner, as a result of which eG agents are able to monitor Java (Servlets, EJB, JSPs) and other non-Java implementations (ASP, PHP, CGI, etc.) with equal felicity.

The eG Web Adapter

The eG web adapter is a layer that fits between the TCP/IP stack and a web server. It can be thought of as a passive probe that watches the requests received by the server and the responses produced by the server. By applying a fast, pattern-matching algorithm on the packets that flow by, the web adapter collects and reports a variety of statistics regarding web sites and the transactions executed by users at these sites without performance degradation. With the help of its web adapter, eG can offer 100% real-time monitoring of 100% of the real transactions (not emulated!), without the need for explicit, expensive logging. Since it monitors all transactions to a web site, this unique capability allows ASPlite to accurately quantify any performance degradations and to proactively alert operators about potential problems with their infrastructure.

Real-time Transaction Monitoring

This section takes the help of an example to illustrate how the eG performs real-time transaction monitoring using the web adapter. Figure 1 depicts this capability.

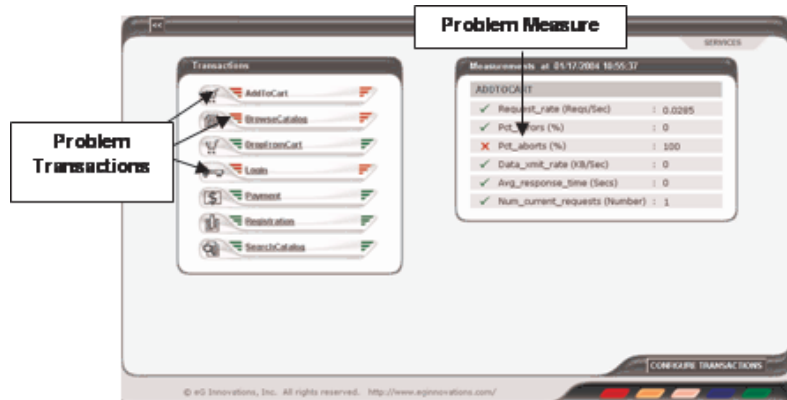


Figure 1 : Real-time monitoring of real web transactions

The scenario shown in Figure 1 demonstrates monitoring of an on-line retail site with the help of eG ASPlite. Users can configure the critical transactions for a site by providing a name for the transaction and the list of pages associated with the transaction. The transactions that have been configured by the user for the retail site in our example are:

1. Add to Cart
2. Browse Catalog
3. Drop from Cart
4. Log-in
5. Registration
6. Search Catalog

The left hand panel of Figure 1 displays color-coded alarms if any of the transactions are not performing effectively. For any chosen transaction, the right hand panel displays the measurements last made by the eG web adapter. The measurements associated with the "AddToCart" transaction chosen in Figure 1 are:

- **Request_rate:** Rate of requests for a specific transaction;
- **Pct_errors:** Percentage of error-filled responses from the web site for a specific transaction;
- **Pct_aborts:** Percentage of requests aborted by users when accessing a specific transaction from a web site (a request is deemed to have been aborted if either the connection is closed without any response being generated, or if the TCP connection is closed as the server is reading the request or as it is responding to the request);
- **Data_xmit_rate:** Rate at which the data is transmitted by the web site in response to user requests for a specific transaction;
- **Avg_response_time:** The average time taken by the web site to respond to requests for a specific transaction, measured in seconds. Only requests for which successful responses are received are considered while computing the average response time.
- **Num_current_requests:** Number of server threads/processes currently in use for serving requests for the specific transaction supported by a web site;



eG ASPlite always pinpoints the reason(s) for why a transaction is not performing to expectation. In the scenario of Figure 1, the failure of the “AddToCart” transaction is associated with a 100% of requests being aborted by users. For further diagnosis, the user can choose any of the problem transactions and can view the topology of the web site (see), which indicates the different application components used by the web site and their individual states.




Figure 2 : Topology view of a web site depicting the problem component(s)

From the color-coding in , it is clear that both the web server and the Netscape application server (NAS server) components are having problems. The direction of the dependency between these components indicates that the problem probably originates in the NAS server component.

Figure 3 : Measurements pertaining to the selected test

Figure 3 depicts the details that eG ASPlite provides for the NAS server component. Each component monitored by eG ASPlite is modeled as comprising of a set of hierarchical layers. The layers of a component are designed such that the state of a layer is impacted by the states of all the layers below it.

For the problem scenario being considered, indicates that the APP_PROCESSES layer of the NAS server component is causing a problem. The state of a layer is dependent on the list of tests that execute on that layer. Now, if you click on the APP_PROCESSES layer, you will find that the ProcessTest of the KJS process associated with that layer is causing the problem. The state of a test depends upon the state of the measures that are collected by that test. In our example, clicking on the ProcessTest (of the KJS process) will reveal the root-cause of the problem with the NAS server – i.e none of the instances of the KJS process are running (The Num_procs_running measure for the KJS process is 0).

Performance of the eG Web Adapter

This section quantifies the performance benefits of the eG web adapter technology. The example under consideration compares the request processing rate of an Apache web server under the following situations:

- A web server that has the web adapter capability enabled for monitoring
- A web server that has transaction logging turned on, and
- A web server that is enabled both with the web adapter and logging

The graph of Figure 4 depicts the variations in the request processing rate of the web server as the input load increases, for each of these situations.

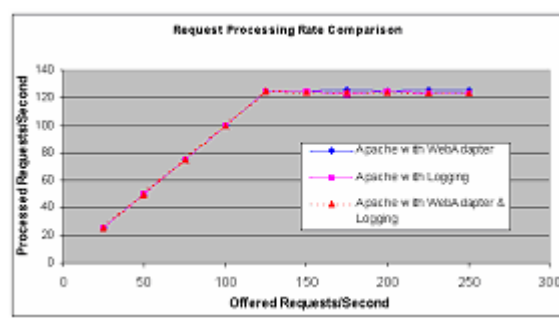


Figure 4 : The performance of an Apache web server with and without the eG web adapter

As can be seen from Figure 4, in an under-loaded scenario, there is no difference in the processing capabilities in the three scenarios considered. The server reaches its capacity at 120 requests/sec. Beyond this limit, the server's processing rate remains almost the same, although the input load is increased. Of the three scenarios, the performance with the web adapter alone is the best. Figure 5 below depicts the percentage of errors seen by requests to the web server as the input load increases. This figure illustrates that at extreme loads, the web adapter offers a close to 2% improvement in performance compared to the use of logging. This experiment does not take into account the overhead involved in processing the log files. For a production web server that receives millions of hits a day, if the overheads of log file processing are considered, the performance improvements are likely to be significant.

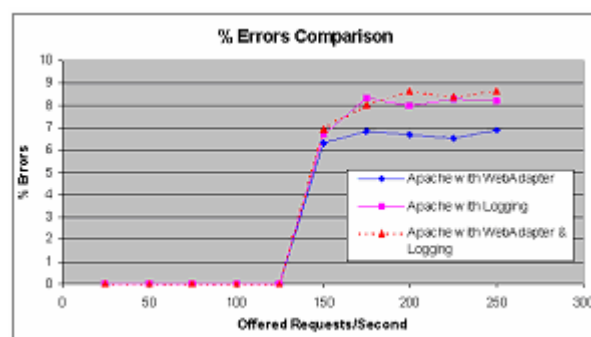


Figure 5 : The performance of an Apache web server as indicated by the percentage of errors generated as the input load is varied for different scenarios

Platforms Supported

The table below lists the platforms supported by the eG web adapter.

OS	Web Servers
Win NT, Windows 2000, Windows 2003	Microsoft IIS, Apache
Solaris	SunONE/iPlanet, Apache
Linux / HP-UX / AIX	Apache, IBM

About eG Innovations

eG Innovations is the leading provider of enterprise-class monitoring and management solutions for IT Infrastructure. The company's 100% web-based monitoring solutions are especially suited for mission-critical infrastructures where proactive monitoring, rapid diagnosis, and instant recovery are critical. Customers worldwide use the eG solutions to improve the quality of their services thus increasing their competitive positioning, lowering their operational costs, and optimizing the usage of their infrastructures.

For More Information

Web : www.eginnovations.com

Email : info@eginnovations.com