

Newport Group, Inc.

IT Trends Research and Reporting

MANAGING WEB APPLICATION PERFORMANCE BEYOND 2000

*New Tools and Services from Mercury Interactive Present
E-business-Driven Customers with More Choice for Achieving and
Managing the Quality of the eCustomer Experience*

INTRODUCTION

Aside from being the driving force behind the Nasdaq's 85.6 percent gain on Wall Street in 1999, the Internet Economy is reshaping how business is conducted. At the center of determining survival in this new and rapidly changing economy, is the continuous availability and reliability of the technologies that make certain the Internet Economy continues to grow both old and new-line businesses. The achievements of e-business leaders like Amazon.com and Charles Schwab continue to inspire others to put the "e" into their own business. Others like Reebok, who retracted from selling product online, or Toys-R-Us who failed to fulfill the majority of 1999 online holiday orders, have struggled. So, why is success possible for some, while others collapse under the weight of the e-business initiative? In a few words, the customer experience. As the primary tool for customer interaction, the Web application plays a significant role in the success of any e-business effort. Protecting the quality and security of the front-facing Web application and its integration with other key enterprise systems by employing continuous testing strategies works to deliver a positive customer experience. Delivering a positive customer experience is a critical element for surviving (and thriving) in the Internet Economy.

This Newport Group white paper looks at several key trends emerging in the area of Web application performance management. Specifically, we will discuss how the use of application performance monitoring is extending testing efforts into the production environment and strengthening the value of existing test assets. Incorporating performance testing in the early design phases of the application development cycle and through deployment is a proven value. Extending that value with application monitoring works to maintain greater control over knowing what users experience in live environments. This white paper will also discuss the growing market for vendor-hosted application testing and monitoring services. Lastly, we will take a closer look at Mercury Interactive and their customer-driven approach to delivering new dimensions in Web performance management, including a production application monitoring tool (Topaz™), a hosted application monitoring service (ActiveWatch™) and a hosted "Test on Demand" service (ActiveTest™) for performance testing.

APPLICATION WATCH: EXTENDING THE VALUE OF AUTOMATED TESTING

In addition to the disaster of downtime, a poorly performing Web application can drain the lifeblood out of a business trying to generate revenue and/or offer enhanced customer service options online. That is why businesses fully entrenched in electronic commerce recognize that it is not only necessary to test the reliability and performance of applications iteratively prior to going live, but that it is also equally important to continue testing efforts into the production environment by employing Web application monitoring strategies and tools. As an early-stage trend, the inclusion of application monitoring functionality into portfolios of testing tools is

aimed at managing Web application performance from the perspective of the end user. For testing tool vendors, application monitoring broadens quality practices throughout the complete application life cycle and promotes collaboration of quality practices among each IT member responsible for the Web application, including design engineers, business analysts, developers, testers and operations engineers.

Pre-deployment test assets, such as scripts and performance result metrics, can be leveraged in the production environment to monitor live application performance. Once pre-deployment testing is completed to satisfy business requirements, the performance metrics deemed acceptable then become the baselines, or comparison measurements for the live site to be measured against. Test scripts can be reused to simulate multiple transaction types (business processes) at consistent time intervals within the live environment to gain an ongoing understanding about whether or not an application is meeting its performance objectives. Traditionally, businesses rely on systems monitoring tools to watch the individual components comprising their computing infrastructure. These tools are equally critical for gaining the detailed information necessary to resolve problems, but tend to monitor and provide information in a stovepipe fashion, or specific to individual layers of the Web enterprise (i.e., network, database, application server, etc.). In short, they do not focus on application performance across each of the application tiers to provide a single, all-inclusive view of what the user actually experiences. Users are concerned with how quickly their request can be filled and therefore businesses providing a service online must be continuously mindful of how users experience their application(s). Application performance monitoring is the high level view - the lookout - signaling alerts to application distress across a growing global population of users.

SURVEILLANCE TACTICS: THE WHY AND HOW OF APPLICATION MONITORING

Web application monitoring tools work to send up red flags when the application under surveillance fails to meet its performance objectives. E-businesses rely on application monitoring tools because their user audience is external and largely unknown. Whereas once, internal audiences of employees may have charged into their IT labs to complain about substandard performance, Web users typically don't bother. They simply take their business elsewhere. So, if Web users are to be kept satisfied, application performance must be maintained for consistency with application monitoring practices.

To know what users are actually experiencing, surveillance tactics such as transaction simulation – the collecting of round-trip response time metrics for transactions conducted by a synthetic user – can gauge whether or not transactions are completing according to business requirements and serve as an immediate alert to problems. It is important to measure all online transaction types (i.e., logon, browse/search, purchase, account query, file download, etc.) on a continuous basis. This is accomplished by driving simulated transactions via the Internet to the application at specific time intervals (i.e., every 15 minutes), measuring response times, validating the accuracy of the returned page responses and then comparing those response times against baseline performance objectives. Ideally groups of simulated transactions should be driven to the application and measured from multiple geographic locations to gain a balanced performance characterization of the global Web user. If businesses have characterized their Web-user populations, having the ability to monitor transactions from more strategic geographic locations that align with actual user populations will provide more accurate results specific to the needs of a particular business. If abnormalities are discovered, appropriate parties must be alerted and corrective action can begin. In some cases, solving a performance problem may be outside the control of a particular business. However, knowing information about the service being delivered to customers can help in determining a reconfiguration of resources and help businesses stay in control of performance to the extent that it is their responsibility. Employing Web application monitoring allows IT groups to keep their finger on the pulse of the user experience and works in the following way:

Securing Baseline Application Performance in Production: For application monitoring to be effective, businesses must first define their performance objectives, (acceptable response time limits) for each application transaction type. Performance test results — gained by testing an application's performance prior to its live debut — can provide appropriate baselines unique to a specific application environment. For example, if test results show that a typical purchase transaction consistently required 45 seconds to process, it would then become the logical baseline to measure against in the live environment. Once acceptable levels of application performance are established, organizations possess key information from which to draw comparisons. IT departments are thereby empowered to set, support or defend application performance objectives. Important to note is that performance baselines may be further normalized within the live environment to more accurately represent response times based on times of day and/or geographic location.



Knowing What Your Users See: One way to gain a perspective on what Web users see is to become one of them. To do this, create scripted monitoring profiles (or leverage pre-deployment test scripts) to simulate typical transactions and continuously compare live response times against performance baselines. Additionally, the accuracy of the page content of each response to a given request must also be validated. The value of a rapid response time is cancelled out if the information returned is inaccurate.

Time intervals for monitoring each transaction must be defined to determine how often transactions are to be fired off for response time data collection (i.e., every 15 minutes, every hour, etc.). Most importantly, plan for appropriate paths of corrective action in the event that response times fall outside defined baselines. Plan for who needs to be notified and by what method. Anticipate instances where a solution can be invoked automatically, such as a Web server reset. Once this basic information has been established, utilize an application monitoring tool and/or service to activate the plan.

Proactive Alerts to Control Performance: The primary reason for monitoring application performance in the first place is to be alerted immediately to performance abnormalities in the live environment. For example, maybe traffic volumes are increasing dramatically and are stretching system capacity. This could happen as a result of a major company news event, a widespread advertising campaign or business growth. Being alerted to abnormalities immediately provides the opportunity to take corrective action before all users are affected and the issue grows more complex. Base alerting methods on the seriousness of the performance problem and automate corrective actions where appropriate (i.e., re-booting a problematic Web server, re-routing traffic to another resource). True proactive application management begins by capitalizing on automated processes designed to diagnose, alert and resolve common performance issues around the clock.

Strengthening Quality Practices Throughout the Web Application Life Cycle: Response time data collected from the live site becomes useful for testing the performance of later releases of the Web application. Trended monitoring data helps to characterize performance in the production environment by busiest times of day and geographic locations. It is a certainty that Web applications will evolve and change as new technologies emerge and as customer demand requires. Trended response time metrics align business requirements with realistic performance activity to more accurately measure performance improvements or degradations over time as an application is changed and re-released. The goal of a continuous testing process is to closely align off-line performance tests with the behavior of the live environment, thus allowing performance testing to become more accurate over time. It also helps to close the quality loop, making certain that testing becomes a life cycle process for Web applications.

Trended performance data provides a historical perspective on how well customers are serviced over time. Generating service level reports from the trended data is a good way to demonstrate service level efficiency to managers, business partners and when appropriate, venture capitalists to chronicle efforts towards ensuring adequate customer service. Reports can be used as documented proof of whether or not a Web application is meeting business required performance objectives or existing service level agreements.

Collecting the Information to Aid in Problem Resolution: Web application monitoring solutions collect data surrounding transaction response times and validate the accuracy of returned application responses to customer transactions and/or queries, allowing problems to be uncovered and examined. The tools work to isolate where problems are stemming from in terms of time of day, geographic location and specific transaction. Once the problematic transaction is exposed, operations engineers will need to do additional work to uncover what component of the Web infrastructure caused the issue (i.e., is it the client, server, or network? Is it within my control or do I involve my ISP to allocate more resources or bandwidth?). Monitoring tools that allow the further dissection of a transaction to isolate the root cause of an issue work to expedite problem identification and resolution.

TEST ON DEMAND: HOSTED PERFORMANCE TESTING SERVICES

Knowing how to properly design and implement meaningful tests can seem a daunting task in the face of extreme time-to-market pressures. The Internet Economy has elevated the business struggle of balancing time-to-market pressures with the time-intensive task of setting up an effective test environment complete with staff, expertise, hardware and automated testing tools. This struggle has led to another significant trend beginning to develop for software testing. That trend centers on vendors stepping up to the plate to offer hosted “Test on Demand” services. Utilizing their own tools and engineering expertise, vendors are now offering testing as a service to be conducted over the Internet. *Test on Demand* services offer new and viable options to a wide range of customers and by co-locating where services are hosted, vendors can work to deliver services independent of time of day. In general, deliverables from hosted test services include consulting time, a set of tailored test scripts, the application test run and final report providing details about the tests and suggestions for areas in need of improvement. Customers of

the hosted service watch tests happen in real-time via any standard Web-browser to see how their application fairs. Armed with information, they can then initiate changes to make improvements to their applications.

Newport Group believes that for businesses updating applications frequently — as is typical for e-businesses — the most cost effective technique over the long-term (given current pricing structures for hosted services) would be to have testing tools and resources in-house to cover the majority of testing responsibilities. However, the current combination of several market conditions will work to quickly fuel **Test on Demand** services. First, there is the exploding demand for more complex e-business capabilities that integrate seamlessly with existing business systems to deliver revenue and/or customer service value. Development projects are moving beyond simple front-end Web applications to include more complex requirements and integrations. As complexity increases, application performance will remain the top concern for businesses. Second, there is the speed at which companies are demanding the completion of these projects. Third, because these applications are available to many via the Web, the precision of application execution and its ability to scale to handle thousands of concurrent users is paramount to business integrity, and that means extensive testing. Yet, there is a shortage of QA engineers fluent in the use of available tools. Learning how to effectively leverage test automation requires a learning curve that is in direct contradiction with time-to-market pressure. The combination of these market conditions is causing businesses to seek the assistance of outsourced services.

Newport Group recognizes real value in hosted testing services for the following scenarios: (1) for businesses new to test automation, with limited resource bandwidth, looking to get a jump-start on establishing a solid test foundation, with eventual plans to bring tools in-house as their business grows and their application evolves; (2) for businesses equipped with adequate resources, but unable to tap those resources due to project delivery scheduling conflicts; (3) for any business looking to gain a second opinion on the performance health of their application from an objective source outside its own engineering team. This can be especially useful for businesses orchestrating a major business news announcement or advertising campaign that has the potential to drive high volumes of increased site traffic. In this latter scenario, consider the **Test on Demand** service as a rider for the business test insurance policy. If due diligence testing throughout the application life cycle has been performed in-house, businesses can gain an additional level of confidence from “another set of eyes” evaluating their application performance. For companies opting not to bring tools in-house, the benefit of iterative testing from application build to build is either lost or can become expensive via hosted test services.

As this very early trend develops and customer demand for services intensifies, Newport Group expects to see service delivery models and pricing structures evolve and become more flexible. Newport Group also believes some testing tool vendors will adapt the ASP (application service provider) business model for a range of testing solutions. In that respect, customers would maintain control over their business test process by renting testing applications for a monthly service fee based on usage, but would be responsible for designing and developing tests and interpreting subsequent test results.

NEW DIMENSIONS IN WEB PERFORMANCE MANAGEMENT FROM MERCURY INTERACTIVE: TOPAZ™, ACTIVEWATCH™ & ACTIVETEST™

In October of 1999 Mercury Interactive introduced Topaz, a Web application performance management tool designed to provide real time monitoring from the perspective of the end user. Topaz enables the monitoring of round-trip transaction response times by firing simulated user transactions at a live Web application. This process begins by recording application business transactions and then setting them up to be replayed on the production site at pre-defined time intervals. Topaz leverages Active Agents™, a technology used to collect performance data while transactions occur, to monitor and report on the availability of a given application. Topaz also validates the content of page responses delivered to user requests to ensure the accuracy of the information being delivered by the Web application. If performance abnormalities occur, Topaz activates its Alert Wizard™ to send a notification before issues have time to escalate. Topaz is programmable to monitor exactly when and how often a business requires. This allows for application performance scrutiny at known peak demand times. The resulting data is converted into browser-based reports that give any number of people within an organization the chance to view application performance metrics. More specifically, the reports have a “drill down” capability to assist in isolating a particular geographic location or transaction that may be causing a performance issue. Transactions can further be dissected to view the portion of total response time that was spent traversing the network and the portion that was spent in the server. This response time detail aids in isolating performance issues more quickly to in turn



implement a resolution more quickly. Topaz's Graph Gallery™ provides a set of pre-defined HTML reports for viewing at any time, from any browser. The advantage of this feature is centralized access of historical data that can be used for trend analysis.

In addition to offering Topaz to customers directly, Mercury Interactive utilizes Topaz to power its ActiveWatch service. In essence, ActiveWatch is a hosted version of the Topaz monitoring technology. Utilizing multiple, geographically dispersed ISPs throughout the world, businesses can monitor all of their critical application business processes to understand the customer experience. ActiveWatch can send business transactions from over 400 distinct points-of-presence (POPs) around the world. To do so, the ActiveWatch service leverages the combined infrastructures of several Internet hosting companies, including Exodus Communications, UUNET (an MCI WorldCom Company), Qwest Communications, AT&T and GTS.

Different from competing monitoring services that place monitoring agents at the major data centers of Internet hosting companies to drive simulated transactions and collect measurements, Mercury Interactive drives transactions from their main data center in Sunnyvale, CA through specific POPs as defined by the customer. Referred to as Mercury Interactive's **POPs-On-Demand** option, ActiveWatch customers may choose which POPs will most accurately reflect their actual user population. Once the appropriate POPs are defined and chosen, simulated transactions will be driven from Mercury Interactive's operations center and through the select locations. Eventually all transactions will funnel to one of the main data centers that connect to the Internet backbone, but by driving transactions from the Mercury Interactive operations center through specific, local POPs that are then connected to the main data centers of the hosting companies via local access routers, ActiveWatch can get a more accurate read on the *last mile measurements*, in order to gain more precise performance monitoring metrics.

To extend the **POPs-On-Demand** option, Mercury Interactive leverages its agreements with Qwest Communications, AT&T and others to put a dedicated T1 line between the ActiveWatch operations center and the specific, customer-requested locations where local access routers are housed, providing ActiveWatch customers with a much broader choice of POPs to collect measurements from, while simultaneously providing them with a way to monitor their transactions from the locations that are most closely aligned with their Web application user base. Transaction response time results are made available via Internet access, and can also be sent in the form of an aggregated performance report via e-mail on a daily basis. The hosted service eliminates the logistics and hardware purchases for organizations seeking to monitor and understand application performance from the perspective of their global customers.

As well as Mercury Interactive's own initiatives with Topaz and Topaz-powered services like ActiveWatch, Mercury also co-developed analogous technology in partnership with IBM-owned Tivoli Systems. Tivoli in turn has optimized the base technology in their efforts to fold it into their Tivoli Application Performance Management (TAPM) product suite and meet the needs of Tivoli customers. In addition, Mercury Interactive OEM's its ActiveWatch service to other organizations wishing to offer hosted application monitoring services under a private label. As a new tool in a rapidly expanding market, Newport Group expects to

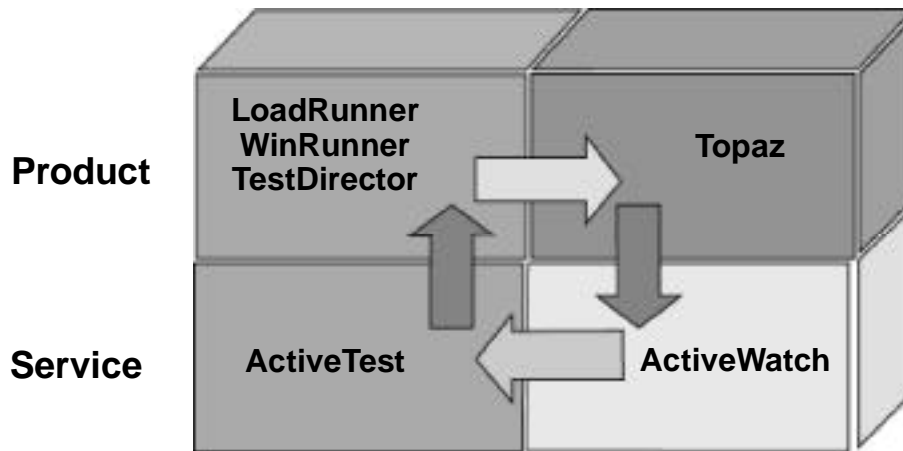
"We adopted ActiveWatch because we were very interested in testing our site from the customer perspective. Lexis-Nexis is highly instrumented internally, so we have all the response time metrics from our perspective, but there was nothing that was being done from a customer perspective. ActiveWatch serves as a sanity check on the metrics collected and reported internally. It completes the application performance picture. With a Web application, the majority of the response time (about 60%) is out of our control and there are times when our TCP/IP stack is handed off and our internal metrics lose track of it . . . meaning we don't know if there is a problem. However, customers will perceive any problem as our responsibility. Therefore, we wanted to get a handle on how our customers perceived our application. To do this we implemented ActiveWatch to monitor our site every 15 minutes from 9 points-of-presence around the world.

The Lexis-Nexis site, which is a research subscription site that receives over 1.7 million hits per day, is very dynamic. There is a lot of activity that happens in our back-end systems as a result of a research query in need of a response. Multiple steps must happen when a customer comes to the site. First they must log in with a secure ID and password, but then they do have the option to stay within the secure part of our application or go outside of it. To do this, our application invokes re-direct URLs and a repainting of the screen. Mercury Interactive was the only vendor we found that could handle the dynamic nature of our application and provide us with meaningful response time metrics. My overall experience in working with them [Mercury Interactive] has been great. They are very responsive and open to our needs for the ActiveWatch service."

-Robert Forrest, Capacity Performance Monitoring Consultant
-Speaking on his use of ActiveWatch
-Lexis-Nexis



Mercury Interactive's Web Performance Management Solution



- *Scripts are 100% reusable from testing to production*
- *Customers can mix & match solutions to meet their needs*

see further product enhancements to Topaz and ActiveWatch that include data collection methods that deliver more granularity for response time metrics and additional options for delivering performance reports and data to customers wishing to correlate ActiveWatch data with internal metrics gained by other methods such as systems monitoring.

TEST ON DEMAND: THE ACTIVE TEST™ SERVICE

Primarily hosted by Internet Service Providers Digital Island, AT&T, Qwest Communications and Loudcloud and driven by Mercury Interactive's LoadRunner®, ActiveTest is an automated testing service designed to verify the performance of Web applications over the Internet. Upon signing on for the ActiveTest service, Mercury Interactive's test engineers work to understand the customer Web application and corresponding performance objectives. Next, Mercury Interactive's test engineers design tests and profile an Internet user population that closely resembles what is anticipated in the live environment. Then, using LoadRunner, a tailored set of test scripts are developed and user loads consisting of hundreds or thousands of simulated transactions are driven against the application. Leveraging the current hosted infrastructure, the ActiveTest service can drive transaction load simulations of 100,000 concurrent users or an approximate Web server load of three billion hits per day.

As load is incrementally increased, both the test engineer and the customer monitor application behavior in real-time via any standard Web browser. ActiveTest measures transaction response times to aid in defining the performance boundaries of an application within its unique environment. By measuring the response times between the client, network and server, ActiveTest monitors can isolate problems within the Web enterprise (is it the firewall, the Web server, the application server, the network, etc.?). ActiveTest results reveal which areas of the application are in need of tuning to achieve better performance. Based on customer observations and summary performance test reports generated by the ActiveTest service, customers can assess their application weaknesses and make improvements to their systems. Further, results define various server capacities, information that is necessary to manage growth (i.e., when is it time to add more headroom to the system).

According to Mercury Interactive, ActiveTest customers are guaranteed a turnaround time of five business days or less. Upon completion of the ActiveTest service, customers retain any and all load testing scripts and performance reports compiled by Mercury Interactive's test engineers. These tests can then be reused for either subsequent ActiveTest runs, or internally when a customer buys LoadRunner. In addition, the same scripts can be used for

“In the Internet space, businesses are constantly pushed to deliver applications quickly. In our case, we used Mercury Interactive’s ActiveTest service to save time and reveal unexpected performance issues surrounding network bandwidth and server configurations, while preparing for our public launch. Not knowing this information prior to launch could have caused significant embarrassment. Using ActiveTest was like watching the heartbeat of our servers as the Mercury Interactive consultant stepped on and off the gas to test varying amounts of virtual users accessing the application. We were able to watch performance monitors over the Internet as our application was tested and we were immediately made aware of its weaknesses. We now know what the maximum concurrent usage is per server, which helps us determine when its time to add more capacity to our system. That’s the beauty of the ActiveTest program.

ActiveTest is a verification service. Having Mercury Interactive testing tools in-house, we are diligent about testing our application internally, but ActiveTest provided validation of what we were doing right. Additionally, we adopted the ActiveWatch service to monitor our application limits; we monitor several transactions every 15 minutes from 8 different geographic locations and log into Mercury Interactive’s Web site once per day to review our numbers. The end result is that we are satisfied with our return on investment for both services because we are extremely mindful of performance and constantly review our numbers to make improvements to every single release. It’s our way of showing our investors and partners that managing application performance is a top priority at Bitlocker.”

-Deanna Falcon, Director of Customer Care and QA
-Speaking on her use of ActiveTest and ActiveWatch
-Bitlocker, Inc.

post-deployment application monitoring by Topaz or the ActiveWatch service. At twice the cost of the average price, Mercury Interactive offers an “Emergency” component to the ActiveTest service, designed to test a Web application and provide performance results within 24 hours.

THE BOTTOM LINE

If there’s one thing that’s certain, it’s that survival in the Internet Economy will be buoyed by Web applications that are able to execute business processes with speed and precision. If businesses are to succeed in building customer reliance on Web systems, they must support application testing and monitoring as a business necessity. With the addition of the LoadRunner-driven ActiveTest service, Topaz, and the ActiveWatch service, Mercury Interactive provides a broad range of choice for customers in need of integrated application testing and monitoring solutions. Expanding on their existing portfolio of test solutions currently serving Web, ERP, client/server and legacy applications, Mercury Interactive’s ActiveTest service provides e-business-driven customers with a new choice for leveraging LoadRunner and gaining a sanity check on their application performance. ActiveTest is aimed at customers needing to balance time-to-market pressures against the lack of test resources and/or in-house expertise that so often causes test initiatives to fail or be shortcut. It can also serve well for businesses looking for the added confidence and objectivity of having someone outside their engineering department test the performance of their application. In addition, Mercury Interactive has made major investments in developing Topaz and the ActiveWatch service to offer solutions that extend pre-deployment testing tasks into the production environment and complete the application life cycle quality loop. Given the nature of the Internet itself, it is true that there are performance issues that fall outside the control of any business. However, new tools and monitoring services from Mercury Interactive do provide businesses with a method to employ surveillance tactics to the application performance issues that are within their control.

Newport Group believes that the Internet Economy will drive significant growth in Web application monitoring tools and services as businesses seek to understand how users continuously experience their applications. Despite the most thorough pre-deployment testing efforts, which do work to shake out the major performance showstoppers prior to releasing a site for general use, there are likely to be unpredictable occurrences that can stress or even crash a Web application once it is open for public use. Incorporating careful fail over and redundancy strategies can help to temper how the users are affected and monitoring live application performance is a practice that can limit the business risk for production mishaps. Being alerted to production problems immediately is a critical management advantage for businesses seeking to deliver consistent application performance. Simply put, it is highly unlikely that the majority of Web users will call to complain about poor online service. Customers and business partners will take their business elsewhere if they become dissatisfied. Therefore, it is important for businesses to seriously consider taking proactive steps to manage application performance before issues grow into serious business problems.



About Newport Group

Newport Group is an independent information technology research firm. Founded in 1997, Newport Group was created with the intent and interest to provide detailed research services to major corporations and software vendors that share a vital interest in information technology. Drawing on years of IT research experience, Newport Group concentrates on selective IT research topics and trends. The benefits of this exclusive research approach are passed on to the client base with demonstrated expertise and insight.

The information contained in any and all final written materials by Newport Group will be based on accurate reliable information. However, because Newport Group information is gathered through numerous sources, including third parties, we cannot warrant that the final publication is complete and error-free. Newport Group disclaims all implied warranties, including, without limitation, warranties of merchantability of fitness for a particular purpose. Newport Group or its affiliates shall have no liability for any direct, incidental, special or consequential damages or lost profits.

© Copyright 2000 Newport Group, Inc. Reproduction is forbidden unless authorized
Additional copies may be obtained by contacting 1-508-362-9799

3230 Main Street, P.O. Box 909, Barnstable, Massachusetts 02630
Telephone: 508-362-9799 www.newport-group-inc.com Fax: 508-362-5229