# Scalable Web Deployment with Oracle Developer Server

*Technical White Paper*
*November 1998*

## A Benchmark Comparison of Client/Server and Web Deployment by Retek Information Systems

ORACLE®
Enabling the Information Age™

Retek

# Scalable Web Deployment with Oracle Developer Server

## A Benchmark Comparison of Client/Server and Web Deployment by Retek Information Systems

### OVERVIEW

Metcalfe's law observes that the value of a network increases dramatically with each node that is attached to it. This simple law means that as the Internet grows and becomes accessible to more people from more locations, the benefit of deploying applications to the Web also increases dramatically. Unfortunately, many organizations have to weigh this benefit against the cost of rewriting their existing client/server applications using a Web-enabled technology such as Java. Oracle Developer Server provides a solution to many of the problems associated with moving to the Internet. With Oracle Developer Server, Oracle Corporation provides a seamless and rapid transition from client/server to the Web.

This paper is the result of an independent research study by Retek Information Systems. It describes the benefits associated with Internet deployment and explains the results of the performance tests, showing comparative results for Oracle Developer Server, Client/Server, and Windows Terminal Server.

### UNDERSTANDING THE BENEFIT

The press and many organizations tend to miss the full value of Web-enabled applications by focusing on a single benefit - lowered total cost of ownership (TCO). Although reduced TCO is an important benefit for Internet deployment, it certainly does not encompass the full power of the Web. To understand the benefit of Internet deployment, the easiest place to start is the driving technologies. That is, to examine the driving factors that create the need for Web-enabled applications. These factors include:

- Low cost of high-speed Internet/intranet access
- Ubiquitous and standardized access to the Internet
- Internet Open standards

The low cost of high-speed Internet/intranet access brought about by the proliferation of communication technologies, such as cable modems and ADSL, enables organizations to cost effectively push information and functionality to a large number and wide variety of users. In the retail sector, for example, organizations are now beginning to push more transactional and data warehouse information to stores and warehouses, and to enable functionality at these locations, such as ordering. Furthermore, some retail organizations are realizing that they should look beyond the traditional user base and are pushing information to partners as well. For example, information such as supplier performance and product availability is of value to the supplier, the store operator, and executive management. Some organizations have even taken the concept a step farther and are "virtually extending the store." This is being accomplished by using Internet-enabled kiosks to run Web-enabled customer ordering applications in the store and allowing customers to purchase large items, such as appliances, that would have been impossible to stock in the store previously. All of this is possible with the advent of inexpensive and high-speed bandwidth that now allows stores, warehouses, and other remote locations to be connected to Web-enabled information and applications at a reasonable cost.

While inexpensive, high-speed connections make it cost effective to distribute information and applications to fixed locations, ubiquitous and standardized access to the Internet has a huge impact on the number and variety of users that information and applications can reach. The rapid expansion of Internet availability throughout the world today connects many businesses, homes, airports, and soon even hotel rooms and other locations. Because of this availability, information and applications can now be cost effectively provided to new types of users. In the retail sector, for example, remote buyers on buying trips can dial in and retrieve open to buy information and previous order information. Furthermore they can enter new orders and collaborate with supervisors by transmitting images and other important information.

Some organizations are finding that the information previously held only within their corporate databases is extremely useful in the hands of customers. As more and more customers have access to the Internet, not only is selling to them at home an option, it is also possible to make new product and availability information directly available to them. It is interesting to note that it is actually this new ability to reach a large number of new users that drives the need for low total cost of ownership. As organizations disseminate their information and applications across a large user population, the cost of supporting that total user population must be kept within reasonable boundaries. This is where technology that enables self-service access to information and applications, such as Java, becomes critical to keeping costs within reason.

While the driving factors that are connecting more users are extremely important, the open standards of the Internet remains one of the most overlooked and yet most powerful advantages of Web-enabled applications. There are many benefits in utilizing these open standards, but one of the easiest benefits to realize is integration with other Web applications. On the Web, integrating other Web applications is as easy as pointing an existing application to another Web page. If the integration point changes, the system simply points to a new Web page. With this technology, products can now offer integrated, customized interfaces for users, with very little investment. In a retail example, a Store Managers Workbench Web page could quickly be created that provides a single start point for all store managers in the organization. This Web page could point store managers at a Web-enabled purchase order, receiving and shipping functionality in a merchandise management product. Furthermore, the Web page could point the user at financial information for the store within a Web-enabled financials package. Finally, the Web page could point the store manager at competitors' Web pages to get the latest promotion information. All of this could be encapsulated in a single Web page that store managers can use to launch into the tasks they should be performing on a day-to-day basis.

To summarize, low cost of high-speed Internet/intranet access, ubiquitous and standardized access to the Internet, and Internet Open standards, are all important benefits of Internet deployment.

### UNDERSTANDING THE TECHNOLOGY

Once the benefit of Web-enabled applications is fully understood, it becomes critical to understand the driving technology. In the case of Oracle Developer Server, this means understanding how this product performs against other thin-client solutions, who is using this technology in production, and what the technical considerations are.

To answer these questions, Retek Information Systems conducted an exhaustive set of tests in order to understand and compare the Oracle Developer Server, Windows Terminal Server/Citrix MetaFrame Server, and standard client/server applications. The goal of these tests was to provide a comprehensive understanding of the technical issues and the performance of each platform.

### UNDERSTANDING WIDE AREA NETWORK PERFORMANCE

When data is transmitted over a WAN, performance is gauged in relation to performance over a LAN. The key to understanding the impact of deployment over a WAN is to clearly identify and measure the factors that impact WAN performance. In a properly configured network where factors such as improperly configured equipment are not an issue, bandwidth restraints and packet latency are the two driving factors that affect overall performance.

### BYTES AND BANDWIDTH

A byte represents a piece of information that is sent across a network. On a typical LAN it is theoretically possible to send up to 1,250,000 bytes/second. If you think of a network as a pipe through which information flows, that means that up to 1,250,000 bytes can flow through that pipe in a given second. If users attempt to force more than 1,250,000 bytes of data through the pipe, then the pipe backs up and a delay is experienced. This is typically not a large problem for LAN users because a high amount of data can be pushed through a LAN in any given second.

In a WAN, however, the size of the pipe shrinks drastically. For example, a 56K connection accommodates less than 7,000 bytes/second of network traffic. This means that this smaller pipe is now much easier to fill up and more likely to cause delays. Since a large number of bytes flowing through a small amount of bandwidth can cause bottlenecks, the total number of bytes sent and received is a key performance indicator. If a large number of bytes are being pushed through the network, the pipe will back up and performance will be slow.

### PACKETS AND LATENCY

Bytes are grouped together and sent across the network in a unit called a packet. Another key factor that affects performance is the time it takes a single packet to leave the host and reach the destination. This factor is known as packet latency.

Packet latency can easily be determined with the use of a common network function called ping. Ping sends a single packet across the network and back and measures the time it takes for the round trip. Packet latency is half the time of a ping. If a ping takes .1 seconds, the latency is .05 seconds.

In a typical LAN situation, the packet latency is effectively zero seconds. Under a WAN, this value varies depending on the type of network that is being used. A packet latency of 50 - 60ms is considered relatively good.

On most applications designed for a LAN, a typical conversation between a client and a server machine proceeds like this:

Client: User typed the character 1

Server: Acknowledged

Client: User typed the character 0

Server: Acknowledged

Client: User typed the character 9

Server: Acknowledged

In this type of conversation, network latency can add up quickly. In a simple example, if the latency were 1 second, the additional time due to latency would add up like this:

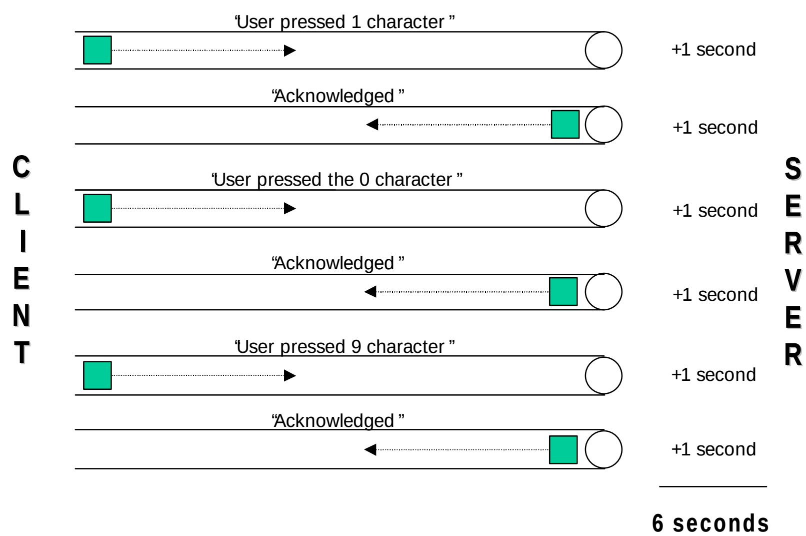


This means that a 6-second latency would occur over this network for this particular transaction. Although packet latency will typically be closer to the 50-100 millisecond range, the total time due to latency can still add up quickly if hundreds or thousands of packets are being sent, which is typical in an application designed to run in a client/server environment.

Note that this latency is not experienced for every packet that is sent over the network. It is essentially only experienced for each packet stream, or group of packets sent in a transaction. The following conversation illustrates the concept of packet streams.

Client: User typed the character 1 (Packet Stream 1)

Client: User typed the character 0 (Packet Stream 1)

Client: User typed the character 9 (Packet Stream 1)

Server: Acknowledged (Packet Stream 2)

Server: Show the following form: «ORDHEAD» (Packet Stream 2)

Server: Show the following form: «ORDFIND» (Packet Stream 2)

Client: Acknowledged (Packet Stream 3)

While the transaction caused 7 packets to be sent, there were only 3 packet streams. With the same packet latency of one second, the total delay due to latency would then be as follows:

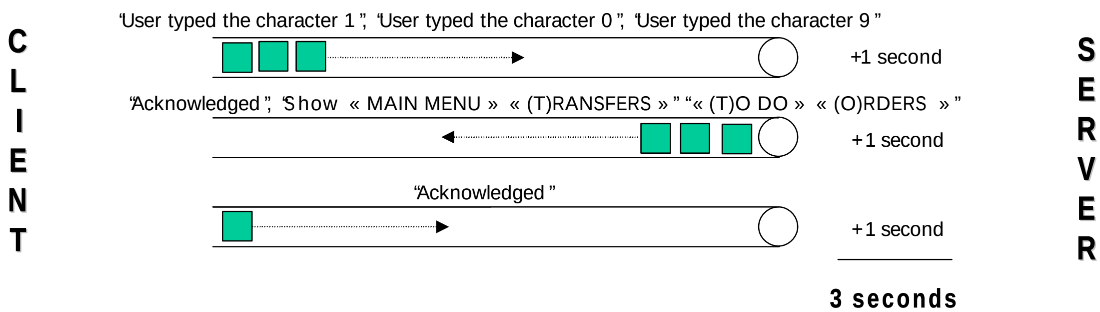


In the above case, even though more packets were sent than in the previous example, only a 3-second latency would occur over this network for this transaction. This is because the packets in the second example were streamed in bundles. So, while the first packet took 1 second to get across to the server, the second packet was right behind, so the server received that packet immediately after the first packet, and so on. This means that the most efficient applications are those with the least number of packet streams.

## OTHER FACTORS

Although bytes and packet streams drive performance, other factors can have a similar impact on performance. Network architectures, network configurations, and protocol architectures all affect overall performance.

For example, in the network architecture category, different types of networks perform differently and introduce variables that are not possible to measure. For example, in a frame relay connection, the intranet line is shared with a large number of other users throughout the world. Given the nature of TCP/IP, a packet will not be placed on the network line until the line is clear. This means that the network will check to see if it can place the packet on the line, and if the line is busy, will wait a predetermined amount of time and then try again. The delay experienced due to this process will varyby the usage and includes traffic that is beyond the known remote users.

## TESTING PERFORMANCE

To understand the overall performance, a series of automated scripts were designed to execute Retek's Oracle Developer applications over the various deployment platforms (client/server, Oracle Developer Server, and Windows Terminal Server/Citrix Metaframe). These scripts were designed based on commonly used Retek functions and included steps that would try to test the weaknesses of each of the environments. The scripts were then executed over a local area network and then over a variety of wide area network architectures. During execution of the test scripts, critical values on the client machine, such as actual time to complete each step, number of bytes transmitted, and number of packets transmitted, were recorded. Furthermore, values such as CPU usage and memory consumption on the server were recorded to understand the overall sizing requirements.

## RESULTS

The tests resulted in a 238-page document outlining the results of each executed step and a summary of the results of the test itself. From that data two simple summary tables were assembled that give an excellent, and surprising, picture of the various deployment platforms:

**Table1: Average and Maximum Bytes/Second Deployment method**

| Deployment Type | Client/Server | Developer Server Java | Terminal Server/MetaFrame |
|---|---|---|---|
| **Average** | 9,453 | 1,572 | 4,410 |
| **Maximum** | 32,855 | 4,239 | 13,107 |

**Table 2: Average and Maximum Packets by Deployment Method**

| Deployment Type | Client/Server | Developer Server Java | Terminal Server/MetaFrame |
|---|---|---|---|
| **Average** | 349 | 31 | 117 |
| **Maximum** | 10,013 | 110 | 1,394 |

These results show that the Oracle Developer Server software generates far fewer bytes and packets than Windows Terminal Server/MetaFrame and client/server. Overall, this means that Oracle Developer Server is far less network-constrained and has much better performance over a wide area network. The actual time that it took to complete each step confirmed this conclusion.

From an application server standpoint, the values were also quite surprising. The following table summarizes overall application server requirements:

**Table 3: Application Server Requirements**

| System Resources Summary | Oracle Developer Server | Windows Terminal Server/MetaFrame |
|---|---|---|
| Average Max User Load (Megabytes) | 17.439 | 24.621 |
| Average %CPU Time | 2.823 | 14.253 |

These results show that the application server for Oracle Developer Server requires far less memory and far less CPU per user. This results in significant cost savings in initial and maintenance hardware fees when deploying a Web-enabled application. In fact, due to these requirements, the actual initial cost to deploy 1,500 name (500 concurrent) Windows Terminal Server/MetaFrame users is estimated at

$1,200,864 while the cost to deploy the same number of users with Oracle Developer Server is $266,286.

Note that the Oracle Developer Server memory values and, therefore, costs are conservatively high because, unlike Windows Terminal Server/MetaFrame, Oracle Developer Server does not require the memory used to be real memory to maximize performance. Due to the nature of memory usage in Oracle Developer Server, a large percentage of this memory can be virtual memory without impacting performance. Real memory required by Oracle Developer Server is actually between 6-10 MB per user. This low real memory requirement means that actual hardware costs should be lower than the amount estimated.

## OTHER TECHNICAL ISSUES

Although performance and cost are critical components in a Web-enabled solution, almost more important are technical questions such as known issues, robustness, and whether this system is production tested and reliable.

### ROBUSTNESS AND KNOWN ISSUES

The Oracle Developer Server has been in development for over a year and half. In that time, the Oracle Tools group has done an excellent job of continuing to enhance the performance and capabilities of the solution. Oracle Developer 1.6.1 and Oracle Developer 2.1, which share almost identical server engines, are the first versions to be truly robust and functional enough to replace client/server. In these releases, almost all technical and GUI issues have been resolved, including the inclusion of shortcut keys (hot keys), extremely reliable server robustness, and ease of installation for both the client and the server.

One of the major benefits of these releases is the availability of Oracle JInitiator. Oracle JInitiator is used to deploy Oracle Developer Server software within web browsers. Retek has been using Oracle JInitiator within Internet Explorer and Netscape with great success. Oracle JInitiator has the following advantages:

- Provides a seamless client deployment vehicle for all major web browsers
- Provides full support for the Oracle-patched JDK which enables enterprise developers to deploy fully-functional Oracle Developer Server solutions
- Java JAR-file caching, which means that startup times over a wide area network are now extremely quick

The biggest issue to be aware of when deploying an Oracle Developer Server application is that Java is CPU intensive on the client. This means that on older machines, the software won't run as fast as client/server on a LAN. At the Pentium II processor level this issue no longer exists, as the processor is able to render Java as fast as client/server.

### PRODUCTION USE OF THE SOFTWARE

As of July 1998, Retek Information Systems currently has one live customer using one of Retek's software packages as an Oracle Developer Server-enabled application. This customer has found the software to be extremely reliable and robust. In late 1998, this customer will expand its user base to serve over 500 concurrent users. Furthermore, Retek has a large number of customers with plans to deploy the Web software in 1999 and an additional group of clients currently evaluating the move for 1999. Internally, Retek has a large percentage of its development staff developing and testing against Oracle Developer Server with great success.

## CONCLUSION

There are many reasons, beyond such benefits as a lowered total cost of ownership, that organizations should be actively considering as they plan Web deploying their applications. These benefits are being driven by the constantly changing technology of the Internet and intranets. By unlocking the huge stores of valuable information currently stored in corporate databases and pushing that information to new types of users, an organization can realize an exponential increase in the value of its information. Furthermore, Web enabling applications opens up the opportunity to create a whole new breed of integrated applications that actually look and feel integrated to the end user.

Oracle has put a great deal of effort into its Oracle Developer Server solution into a robust, complete solution that can be used in a production environment. This effort has paid off in the currently available release, and customers can now be confident that they can production deploy their Oracle Developer applications anywhere that it makes sense. With Internet time driving technology change at an incredible rate and forcing many IS shops to rewrite their applications from scratch, a solution such as the Oracle Developer Server offers a much-needed, quick, and seamless transition to the Web with minimal effort.

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
+1.650.506.7000
Fax +1.650.506.7200
http://www.oracle.com/

Scalable Web Deployment with Oracle Developer Server
White Paper

Author: Jeff Goke, Retek Information Systems;
Rene Saltzherr, Oracle Corporation

Editor: Fergus Griffin, Oracle Corporation,