IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

ORACLE CORPORATION and ORACLE U.S.A. INC.,	)	
Plaintiffs/Counterclaim Defendants,	)	C.A. No. 06-414-SLR
v.	)	**JURY TRIAL DEMANDED**
EPICREALM LICENSING, LP,	)	**PUBLIC VERSION**
Defendant/Counterclaim Plaintiff.	)	

SECOND DECLARATION OF DR. DAVID FINKEL IN SUPPORT OF EPICREALM'S OPPOSITIONS TO ORACLE'S MOTIONS FOR SUMMARY JUDGMENT OF NONINFRINGEMENT AND INVALIDITY

VOLUME V of V

EXHIBITS 47 to 53

Richard L. Horwitz (#2246)
David E. Moore (#3983)
POTTER ANDERSON & CORROON LLP
Hercules Plaza, 6th Floor
1313 N. Market Street
Wilmington, DE 19899
Tel: (302) 984-6000
rhorwitz@potteranderson.com
dmoore@potteranderson.com

Attorneys for Defendant/Counterclaim Plaintiff epicRealm Licensing, LP

OF COUNSEL:

Harry J. Roper
George S. Bosy
Aaron A. Barlow
Patrick L. Patras
David R. Bennett
Paul D. Margolis
Benjamin J. Bradford
Emily C. Johnson
JENNER & BLOCK
330 N. Wabash Avenue
Chicago, IL 60611-7603
Tel: (312) 923-8305

Dated: August 21, 2008
Public Version Dated: August 29, 2008
880482 / 31393 / Oracle

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

ORACLE CORPORATION and ORACLE U.S.A. INC.,	)	
Plaintiffs,	)	C.A. No. 06-cv-414 (SLR)
v.	)	
EPICREALM LICENSING, LP,	)	
Defendant.	)	

SECOND DECLARATION OF DR. DAVID FINKEL IN SUPPORT OF EPICREALM'S OPPOSITIONS TO ORACLE'S MOTIONS FOR SUMMARY JUDGMENT OF NONINFRINGEMENT AND INVALIDITY

I, David Finkel, declare as follows:

1. I am currently the Co-Director of the Computer and Communications Networks program at Worcester Polytechnic Institute. I have been a professor of Computer Science since 1991. I submit this declaration in support of EpicRealm's Oppositions to Oracle's Motions For Summary Judgment of Noninfringement and Invalidity.

2. I have been involved in studying, evaluating, and teaching about computer networks, network software, and the World Wide Web for over 15 years, and have authored numerous articles in the field. Additionally, I have taught undergraduate and graduate courses in the areas of computer networks and computer communications since 1989. Furthermore, I have taught several courses that specifically focus on Internet technology and the World Wide Web.

3. The Accused Oracle Application Server Products literally infringe Claim 1 of the '554 Patent under both parties' constructions. (*See* 2d Finkel Decl. Ex. 1 at 1-20.)

4. The Accused Oracle Application Server Products literally meet the "intercepting" limitation under either party's construction. (*See* 2d Finkel Decl. Exs. 1, 2.)

5. The "Web server" in all of the Accused Oracle Products processes some, but not all requests.

6. All software always goes through its "same iterative process," *i.e.*, it always follows its sequence of instructions. The software code does not change.

7. In Leaf, all requests are immediately forwarded on from the so-called "Web server" to transaction gateways; the "Web server" in Leaf could not process requests itself. Thus, a request could not be "intercepted" at the "Web server" in Leaf because there is no determination of whether the "Web server" or a "page server" should process the request—the

"Web server" in Leaf would never process a request itself. Therefore, Leaf did not disclose "intercepting."

8. The Accused Oracle Application Server Products literally meet the "releasing" limitation under either party's construction because the "Web server" (OHS) is released to process "other requests." (*See* 2d Finkel Decl. Exs. 1, 2.)

Redacted

Redacted

12. The Accused Oracle Application Server Products literally meet the "dispatcher" limitation under both its ordinary meaning and Oracle's construction. (*See* 2d Finkel Decl. Ex. 1, 2.)

13. Mod_oc4j meets Oracle's construction of "dispatcher" because it is a software program that determines which page server should process a dynamic Web page generation request.

14. Source code is written in modular fashion. For example, within one piece of software there can be one section of code that performs the "Web server" functionality and another section of code that performs the "dispatcher" functionality. Therefore, a request is routed from the "Web server" to the "dispatcher" when the request is routed from the code that performs the "Web server" functionality to the code that performs the "dispatcher" functionality.

15. The "Database Connector" is a Dynamically Linked Library (".dll") that runs in the same executable as the IIS 1.0 "Web server."

16. The Accused Oracle Database Products literally infringe claim 1 of the '554 Patent under both parties' constructions. (*See* 2d Finkel Decl. Ex. 3 at 1-21.)

Redacted

Redacted

22. Security attributes and privileges can come from a dynamic Web page request.

23. The Accused Oracle Web Cache Products literally infringe Claim 1 of the '554 Patent under both parties' constructions. (*See* 2d Finkel Decl. Ex. 5 at 1-19.)

24. The "Web server" (Web Cache) processes some but not all requests.

25. Web Cache is released to process other requests under either party's construction of "releasing." When Web Cache routes a request to a "page server," it is released to process

other requests under both parties' constructions in the same two ways—the same two ways that the "Web server" of the Accused Oracle Application Server Products are released.

Redacted

28. Merely having a "page server" process the request does not necessarily mean that the "releasing" limitation is met. For example, a system could be designed to have the "Web server" process a request in parallel with a "page server," *e.g.*, to help ensure a prompt response to a Web client even if the "page server" subsequently became slow or had a problem.

Redacted

31. The Accused Oracle Products also literally infringe the asserted claims that depend from Claim 1 of the '554 Patent. (*See* 2d Finkel Decl. Exs. 1, 3, 5.)

32. The Accused Oracle Products also literally infringe Claims 9 and 10 of the '554 Patent. (*See* 2d Finkel Decl. Exs. 1, 3, 5.)

33. The Accused Oracle Products also literally infringe Claim 11 of the '554 Patent. (*See* 2d Finkel Decl. Ex. 1, 3, 5.)

34. The Accused Oracle Products also literally infringe Claims 2 and 16 of the '335 Patent. (*See* 2d Finkel Decl. Ex. 2, 4, 6.)

35. Oracle provides various manuals to its customers instructing them how to use the Accused Oracle Products in an infringing manner.

36. The Accused Oracle Products infringe the asserted claims under the doctrine of equivalents. The Accused Oracle Products infringe because the differences between the Accused Oracle Products and the asserted claims are insubstantial from the perspective of a person of ordinary skill in the relevant art.

37. Nowhere in the specification does it discuss that "intercepting a request" requires that the Web server distinguish between dynamic and static content. Nor would one skilled in the art understand that the claims of the patents require such a distinction. Oracle Web Cache makes a determination of whether it will process a request itself (*e.g.*, satisfy the request with data stored in its cache) or it will send the request to an application server for processing.

38. The partitioned architecture of the claimed invention involves adding "page servers" which are capable of processing dynamic Web page requests. The "page servers" alleviate the processing burden placed on the "Web server" by processing some or all of the dynamic Web page requests. When a "page server" processes a request it "frees" or "releases" the "Web server" to process other requests because the "Web server" no longer has to expend resources processing that request.

39. In the patented invention, the "dispatcher" can reside on a different machine than the "Web server" or the "dispatcher" can reside on the same machine as the "Web server." Regardless of where the "dispatcher" is located, it performs the same function. When the "dispatcher" resides on the same machine as the "Web server," some of the advantages of the invention will not be enjoyed. Namely, the Web server will not be able to offload the processing of the requests when a request is routed to the "dispatcher." The Web server, however, may still be able to offload the processing of a request when the request is routed from the "dispatcher"

residing on the Web server machine to a "page server." Thus, the purpose of the "dispatcher" is to "dispatch" requests, not to unburden the "Web server."

Redacted

41. The Oracle Web Agent does not make an informed selection of which page server should process the request. When the Oracle Web Agent determines which "page server" (Oracle7 Server) should process a request, it does not make a selection between Oracle7 Servers that are capable of processing the request. Rather, only one Oracle7 server could process the request. Oracle concedes that making an "informed ***selection***" requires, at least, choosing from more than one page server that is capable of processing the request. Thus, the Oracle Web Agent does not make an "informed selection."

42. The Oracle Web Agent also does not make a selection based on information indicating which page server can more efficiently process a request. Because only one Oracle7 server can process a given request, the information used by the Oracle Web Agent does not indicate the relative efficiencies of the Oracle7 servers. Instead, the information only indicates

which Oracle7 server to connect to and how to connect to that server. Thus, the Oracle Web Agent does not use information indicating which page server can more efficiently process a request in deciding which page sever should process a request.

43. The term "dynamic information" refers to information that can only be established during the execution of a program. Dynamic information does not refer to information that can be established prior to the execution of the program. Put another way, dynamic information is of a nature that can only be determined during runtime execution.

44. SQL*Net is remote data access software. Its fundamental purpose is to allow clients and servers to access Oracle databases from remote locations. SQL*Net is not concerned with the management of dynamic Web page requests. In fact, the term "Web page" is never used in the "Understanding SQL*Net 2.2" reference. Load balancing of connection requests has little to do with metric-based load balancing of dynamic Web page requests.

45. SQL*Net routes requests for database connections which are not the same as dynamic Web page requests as recited in all of the asserted claims, which are a type of "work requests." A request for a database connection attempts to establish a connection between a client and a server; the response to a request for a database connection will not be a dynamic Web page. By comparison, a dynamic Web page request typically results in a dynamic Web page in response and cannot establish a connection between a client and a database.

46. When a client is connecting to a multi-threaded database using SQL*Net 2.2 or SQL*Net 2.3, the request for a database connection sent from the client to the database is a wholly different request for a connection than the request originally received by SQL*Net.

47. One way of implementing the selection of the "least called" server is a round-robin algorithm.

48. Oracle 1.0 does not anticipate any of the asserted claims.

49. Claim 4 requires that the "page server" maintain a connection cache to said one or more data sources. Claim 5 requires that the processing step performed by the "page server" also includes the step of logging into one or more data sources. Oracle 1.0 does not anticipate Claims 4 and 5 of the '554 Patent for the additional reason that it cannot maintain a connection cache or log into one or more data sources. According to Oracle, the "page server" in Oracle 1.0 is the Oracle7 server, *i.e.* an Oracle database server. The data sources in Oracle 1.0 are the Oracle7 database. The Oracle7 server does not need to establish a connection to an Oracle7 database in order to retrieve data from the database because they are located on the same machine.

50. Oracle argues that Claim 5 is met because the "dispatcher" (the Oracle Web Agent) logs into the Oracle7 server. However, Claim 5 states that the "processing" step includes logging into one or more data sources and the Oracle Web Agent does not perform the "processing" step of Claim 1. Indeed, Oracle contends that the processing step is performed by the Oracle7 server, not the Oracle Web Agent.

51. Oracle 2.0 does not anticipate any of the asserted claims.

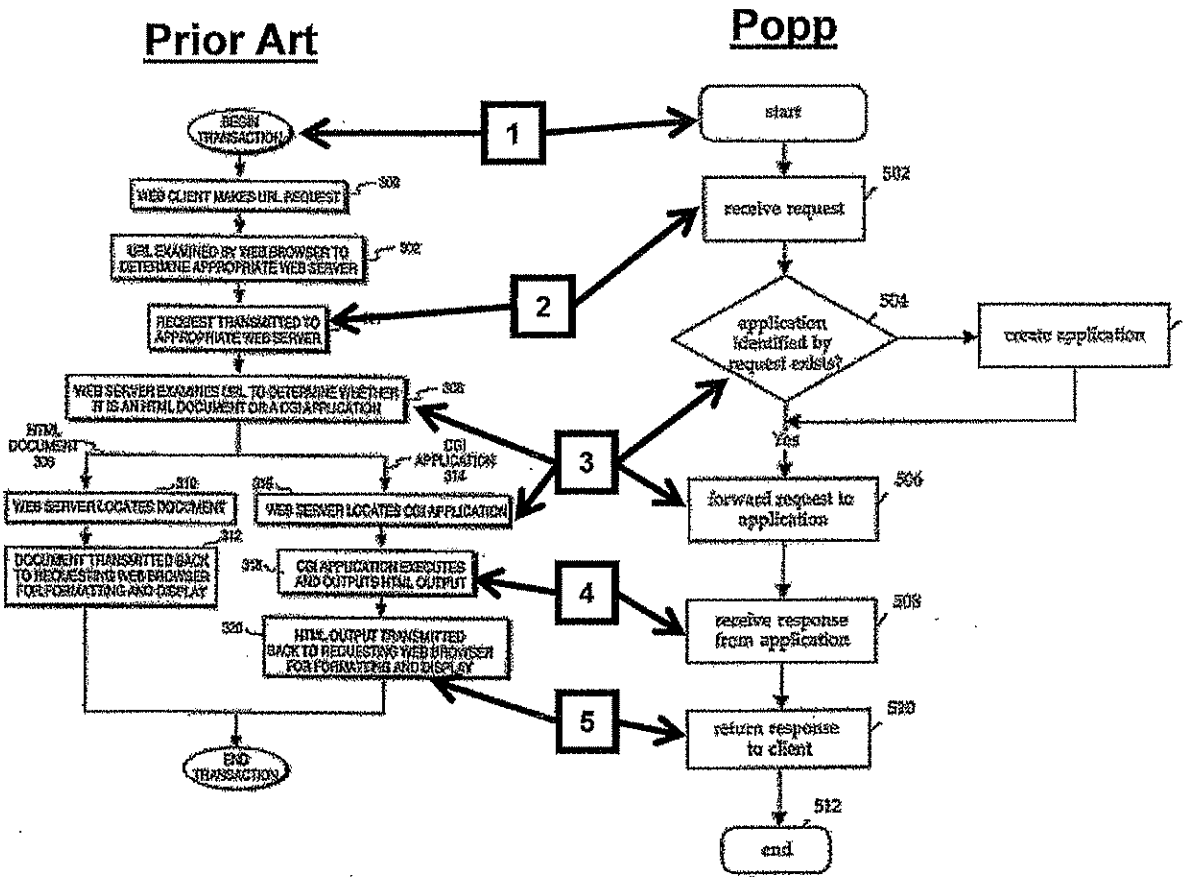
52. Oracle's proposed construction for page server is "page generating software, running on a processor separate from that of the Web server, that generates dynamic Web pages." However, Oracle's alleged "page server" in Oracle 2.0 is a WRBX that runs on the same machine (and therefore the same processor) as the Oracle 2.0 "Web server."

53. Oracle 2.0 does not anticipate Claim 4 for the additional reason that it does not disclose the "page server maintaining a connection cache to said one or more data sources." Under the ordinary meaning of "connection cache," as understood by one skilled in the art, maintaining a connection cache involved a cache of connections to a database that can be

checked out and reused by numerous processes. When a PL/SQL Agent stays connected to a database, there is neither a cache of connections nor can that connection be shared among various processes. Thus, it fails to meet the ordinary meaning of "connection cache."

54. Popp does not anticipate any of the asserted claims.

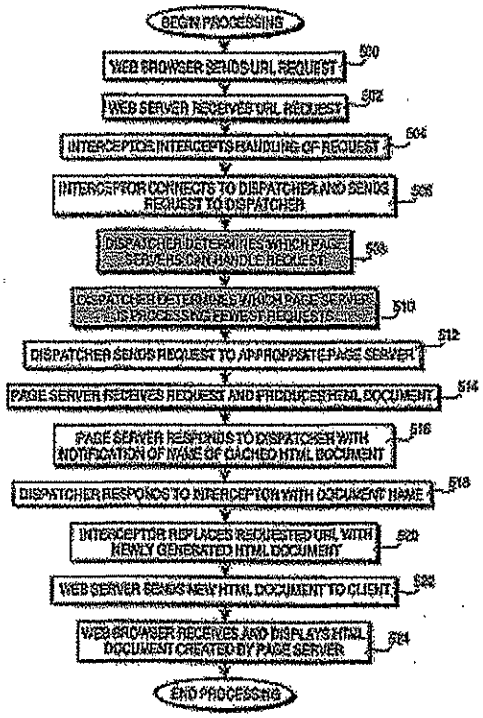
55. Below on the left is a flow chart describing how a request is processed in the prior art Web server disclosed by the '554 Patent and on the right is a flow chart from Popp describing how it processes a request:



nearly identical manner as the prior art disclosed by the patents in suit and considered by the

examiner. Specifically, Popp only determined if the application identified by a request existed, and if it did, forwarded the request to that application.

56. By comparison, in the flow diagram below depicting an embodiment of how the invention of the patents in suit processed a request, the "page server" is selected based on a two-step process (Steps 508 and 510 highlighted in gray below): (1) determine "which page servers can handle request;" and (2) determine "which page server is processing fewest requests" (*i.e.*, which "page server" can more efficiently process the request).



57. Popp's approach to Web site architecture is very similar to Oracle 1.0. In both cases, the primary concern was routing Web page requests to an application that could process the request. Neither Popp, nor Oracle 1.0, gave any consideration to determining which instance of an application "can *more efficiently* process the request." In fact, Oracle 1.0 and Popp are nearly identical to the prior art described in the '554 Patent specification.

58. By comparison, as shown in the embodiment illustrated in Figure 5 of the '554 Patent, the "page server" is selected using steps 508 ("dispatcher determines which page servers can handle request") and 510 ("dispatcher determines which page server is processing fewest requests"), *i.e.*, which "page server" can more efficiently process the request.

59. The CGI interface program in Popp (the alleged "dispatcher") does not make an "informed selection" of which page server should process a request. When the CGI interface program selects an application to process a request, it only determines which application can process a request because there is only one application that can process each request.

60. The CGI interface also does not make a selection based on information indicating which "page server" can more efficiently process a request. Because only one application can process a given request, the information used by the CGI interface does not indicate the relative efficiencies of the applications.

61. The information used by the CGI interface program in Popp is not "dynamic information maintained about page servers." The information relating to which application can process a request and on which server the application is executing is static information because it can be determined before runtime execution.

62. The CGI interface also does not make an "informed selection" because it does not use any information "to make a reasonable or accurate determination as to which [page server] could more efficiently process the request."

63. Dienst does not anticipate any of the asserted claims.

64. Dienst is not concerned with determining which "page server" can more efficiently process a request. Instead, it is concerned with locating the only "page server" that can process a request.

65. Dienst does not disclose "dispatching" for many of the same reasons that Oracle 1.0 and Popp do not disclose "dispatching." All these references are focused on locating the one and only "page server" that can process a given request.

66. The alleged "dispatcher" in Dienst does not make an "*informed selection*" of which page server should process a request. In Dienst, only one "page server" (Dienst Server) contains all the works for a specific publisher. Thus, when the alleged "dispatcher" selects a Dienst Server to process a request, it does not make a selection, but rather locates the only Dienst Server that contains that publisher's works.

67. The alleged "dispatcher" in Dienst also does not make a selection based on information indicating which "page server" can more efficiently process a request. The alleged "dispatcher" only uses information on which Dienst Server is capable of responding to a given request. Because there is only one Dienst Server that can process a request, there is no information on which Dienst Server can more efficiently process that request.

68. The information used by the CGI interface program is not dynamic information maintained about "page servers" indicating which "page server" can more efficiently process a request. Each Dienst Server downloads a listing of the locations of the other Dienst Servers so that it can search the other servers for works. This information is not dynamic information maintained about "page servers" indicating which "page server" can more efficiently process a request.

69. In my opinion, a person of ordinary skill in the art would have a Bachelor's degree in computer science or equivalent, with three to five years of technical experience in software implementation, or equivalent.

70. There is no "dispatching" under either party's proposed construction because the job of the alleged "dispatcher" is limited to locating the one "page server" that can process the given request.

71. Garland does not cure the "dispatching" deficiency in Oracle's invalidity arguments. The "Dispatcher Server" in Garland does not "examine" or "analyze" the request, as required by both parties' "dispatching" constructions. The "Dispatcher Server" also does not send the request to the selected server as required by both parties' "dispatching" constructions.

72. None of the asserted claims are obvious over the combinations of references argued by Oracle.

73. Oracle has not proved by clear and convincing evidence that Garland in combination with any of Oracle's four other references would have rendered obvious to one of ordinary skill in the art any of the twelve asserted claims at the time of the respective inventions.

74. Because Dienst is concerned with providing access to a distributed digital library over the Web, it does not disclose many of the limitations added by the dependent claims of the '554 Patent, which are focused on features to further improve the performance of a Web site. Oracle argues that it would be obvious to combine Dienst with the teachings of Oracle 1.0 or Oracle 2.0 to meet these claims. In addition, none of these references discloses "dispatching." Their combination would not change that fact. Not only does Oracle fail to set forth its basis for claiming that it would be obvious to combine these references, but the secondary considerations argue against such combinations.

75. Further, it would not be obvious to combine Dienst, a reference on providing access to a distributed digital library with Oracle 1.0 or Oracle 2.0. In fact, it would be impossible to use the "page servers" of Oracle 1.0 or Oracle 2.0 to implement Dienst.

76. The information used by JDBC to identify connections that are authenticated with the proper user credentials, user name and password and to identify connections that have the necessary striping can come from the request.

77. I am over the age of eighteen, and, if called to do so, I am prepared and competent to testify on the matters presented in this declaration.

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct. Executed on August 21, 2008, in Texarkana, Texas.

Dr. David Finkel

AUG-21-2008 23:04 P.17

TOTAL P.17

TABLE OF CONTENTS

Exh.	*Document*
1	Claim Chart U.S. Patent No. 5,894,554 – The Accused Oracle Application Server Products
2	Claim Chart U.S. Patent No. 6,415,335 – The Accused Oracle Application Server Products
3	Claim Chart U.S. Patent No. 5,894,554 – The Accused Oracle Database Products
4	Claim Chart U.S. Patent No. 6,415,335 – The Accused Oracle Database Products
5	Claim Chart U.S. Patent 5,894,554 – Accused Oracle Web Cache Products
6	Claim Chart U.S. Patent No. 6,415,335 – The Accused Oracle Web Cache Products
7	Oracle WebServer User's Guide Release 1.0 Part Number A349862" (153)
8	Popp et al. U.S. Patent No. 6,249,291 (195)
9	Oracle WebServer User's Guide Release 2.0.2 Production, Part No. A23646-2 (604/605)
10	Oracle WebServer User's Guide Release 2.0 Production, Part No. A23646-1 and Oracle WebServer 2.0 Technical Note
11	Functional Specification for Runtime Load Balancing in OCI Session Pools, 11gR1, 12/2/05, CHID00081-00105
12	Excerpt from deposition transcript of Lakshminarayanan Chidambaran, Feb. 29, 2008
13	Excerpt from deposition transcript of Carol Colrain, April 15, 2008
14	Excerpt from deposition transcript of Hal Hildebrand, Dec. 21, 2007
15	Excerpt from deposition transcript of Matthew Mayerson, April 15, 2008
16	Excerpt from deposition transcript of Mark Nelson, Feb. 22, 2008
17	Excerpt from deposition transcript of Douglas Surber, Feb. 28, 2008
18	Design Specification for OCI Runtime Connection Load Balancing in 11g, 12/14/05, ORCL01310468-01310498
19	Functional Specification for OFCGI, OHS, iAS, 11AS, 6/2/06, ORCL01209371-01209387

20 Functional Specification for Runtime Load Balancing in OCI Session Pools, 11gR1, 12/2/05, ORCL01166752-01166776

21 Functional Specification for Runtime Load Balancing, RDBMS, 10gR2, 4/1/04, ORCL01060148-01060193

22 Functional/Design Specification for Runtime Work Request Load Balancing using RAC Service Metrics, JDBC, 10gR2, 5/1/04, ORCL01054844-01054858

23 Excerpt from Oracle Application Server Containers for J2EE Support for JavaServer Pages Developer's Guide 10g (9.0.4), B10320-01 Sept. 2003, ORCL00337734-00338101

24 Excerpt from Oracle Application Server Containers for J2EE, Support for JavaServer Pages Developer's Guide 10g Release 2 (10.1.2) B14014-02, July 2005, ORCL00444343-00444558

25 Excerpt from Oracle Application Server Web Cache Administrator's Guide 10g Release 2 (10.1.2) B14045-01, Dec. 2004, ORCL00473184-00473739

26 Excerpt from Oracle Call Interface Programmer's Guide, 11g Release 1 (11.1), B28395-02, Oct. 2007

27 Excerpt from Oracle Call Interface Programmer's Guide, 11g Release 1 (11.1), B28395-03, May 2008

28 Excerpt from Oracle Database Administrator's Guide, 10g Release 2 (10.2), B14231-02, May 2006, ORCL01969895-01970763

29 Excerpt from Oracle HTTP Server Administrator's Guide 10g (9.0.4), B10381-02, March 2004, ORCL00372061-00372362

30 Oracle Real Application Clusters 11g: An Oracle Technical White Paper, April 2007, ORCL01972320-01972333

31 Excerpt from Oracle Web Cache Administration and Deployment Guide, Release 1.0.2, A86722-01, Nov. 21, 2000, ORCL00107955-00108126

32 Excerpt from Oracle XML DB Developer's Guide 11g Release 1 (11.1), October 2007, ORCL01972631-01973585

33 Excerpt from Oracle9i Application Serer: Availability, Scalability and Manageability of J2EE and Web Cache Clusters, An Oracle White Paper, May 2002 ORCL00013356-00013398

34 Excerpt from Oracle9i Application Server Performance Guide, Release 2 (9.0.2), April 2002, A95102-02, ORCL00847081-00847294

35 Excerpt from Oracle9i Application Server, An Oracle White Paper, Sept. 2002, ORCL00010730-00010792

36 Excerpt from Oracle9i Application Server, Best Practices, Release 1 (v1.0.2.2) Sept. 2001 A95201-01, ORCL00190234-00190397

37 Oracle9i Application Server, ORCL00110683

38 Excerpt from Oracle9i JDBC Developer's Guide and Reference Release 2 (9.2) March 2002, ORCL00306997-00307580

39 Excerpt from Oracle9iAS Web Cache Administration and Deployment Guide, Release 2 (9.0.2), May 2002, Part No. A95404-02, ORCL00272342-00272887

40 Excerpt from Oracle9iAX Web Cache Administration and Deployment Guide, Release 2.0.0, Oct. 10, 2001, A90372-04, ORCL00189864-00190165

41 Presentation, Mod_oc4j Load Balancing in iAS 9.0.4, ORCL00018115-00018150

42 Excerpt from Plaintiff's Responses to Defendants' First Set of Interrogatories (1-6)

43 Tanenbaum, A.S., *Computer Networks*, Fourth Edition. Prentice-Hall PTR, Upper Saddle River, NJ, 2003.

44 Oracle WebServer User's Guide Release 1.0, A349862 ORCL00941713-00941960 (MAC003037-003284) (*Ref 153*)

45 US 6,651,108, Popp et al., MAC003907-003927 (*Ref. 194*)

46 US 6,249,291, Popp et al., MAC003928-003968 (*Ref. 195*)

47 Oracle WebServer User's Guide Release 2.0.2 Production, Part No. A23646-2, 1996, ORCL00000782-00001035 (*Ref. 604*)

48 Oracle WebServer 2.0 Technical Note, March 1996, ORCL00000761-00000781 (*Ref. 605*)

49 Oracle and the Internet, Using Oracle in the Emerging Electronic Infrastructure of Business and Society, July 1995, ORCL00937779-00937806

50 Excerpt from Understanding SQL*Net, Release 2.3, A42484-1, 1996, ORCL01808597-01808820

51 The Rush to the Net, May 1996, ORCL01606514-01606526

52 Web Request Broker API, Ver. 0.8, Feb. 12, 1996, ORCL01606852-01606867

53 Excerpt from deposition transcript of Katrina Montinola, Feb. 19, 2008

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

CERTIFICATE OF SERVICE

I, David E. Moore, hereby certify that on August 29, 2008, the attached document was electronically filed with the Clerk of the Court using CM/ECF which will send notification to the registered attorney(s) of record that the document has been filed and is available for viewing and downloading.

I further certify that on August 29, 2008, the attached document was Electronically Mailed to the following person(s):

Mary B. Graham
Morris, Nichols, Arsht & Tunnell LLP
1201 North Market Street
P.O. Box 1347
Wilmington, DE 19899
MGraham@MNAT.com

James G. Gilliland
Igor Shoiket
Townsend and Townsend and Crew LLP
Two Embarcadero Center
Eighth Floor
San Francisco, CA 94111-3834
OracleEpicrealm@townsend.com

Theodore T. Herhold
Robert J. Artuz
Eric M. Hutchins
Eric A. Mercer
Joseph A. Greco
Nitin Gupta
Townsend and Townsend and Crew LLP
379 Lytton Avenue
Palo Alto, CA 94301
OracleEpicrealm@townsend.com

Chad E. King
Townsend and Townsend and Crew LLP
1200 Seventeenth Street
Suite 2700
Denver, CO 80202
OracleEpicrealm@townsend.com

/s/ David E. Moore
Richard L. Horwitz
David E. Moore
Potter Anderson & Corroon LLP
Hercules Plaza – Sixth Floor
1313 North Market Street
P.O. Box 951
Wilmington, DE 19899-0951
(302) 984-6000
rhorwitz@potteranderson.com
dmoore@potteranderson.com

Exhibit 47

Oracle WebServer™ User's Guide

Release 2.0.2 Production

Part No. A23646-2

ORACLE®

Enabling the Information Age

Oracle WebServer User's Guide, 2.0.2

Part No. A23646-2

Copyright © Oracle Corporation 1996

All rights reserved. Printed in the U.S.A.

If you have not read this copyright page, you should read it *in its entirety*. If you have read this page, you can go directly to the Table of Contents.

If you find any errors, omissions, or have any suggestions on how the information in this manual can be improved, please e-mail Oracle WebServer Documentation.

Primary Authors: Martin Gruber, Kennan Rossi

Contributors: Seshu Adunuthula, Mala Anand, Matt Bookman, Walter Hudson, Kelly Ireland, Magnus Lonnroth, Raymond Ng, Mary Owen, Robert Pang, Charles Prael, Ankur Sharma, John Zussman

This software was not developed for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It is the customer's responsibility to take all appropriate measures to ensure the safe use of such applications if the programs are used for such purposes.

This software/documentation contains proprietary information of Oracle Corporation; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited.

If this software/documentation is delivered to a U.S. Government Agency of the Department of Defense, then it is delivered with Restricted Rights and the following legend is applicable:

Restricted Rights Legend Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of DFARS 252.227-7013, Rights in Technical Data and Computer Software (October 1988).

Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

If this software/documentation is delivered to a U.S. Government Agency not within the Department of Defense, then it is delivered with "Restricted Rights", as defined in FAR 52.227-14, Rights in Data - General, including Alternate III (June 1987).

The information in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Oracle Corporation does not warrant that this document is error-free.

Adobe and Acrobat are trademarks of Adobe Systems Incorporated.
Java is a trademark of Sun Microsystems Incorporated

Netscape Navigator is a trademark of Netscape Corporation.
Oracle, SQL*Forms, SQL*DBA, SQL*Loader, SQL*Net andSQL*Plus are registered trademarks of Oracle Corporation.
PL/SQL, Oracle7, Web Request Broker, LiveHTML, Web Access Manager, Oracle Browser, Oracle WebServer Option, Oracle WebServer, Web Agent, Web Desktop, and Web Listener are trademarks of Oracle Corporation.

NOTICE

Copyright 1995 by: Massachusetts Institute of Technology (MIT), INRIA.

This W3C software is being provided by the copyright holders under the following license. By obtaining, using and/or copying this software, you agree that you have read, understand, and will comply with the following terms and conditions.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee or royalty is hereby granted, provided that the full text of this NOTICE appears on ALL copies of the software and documentation or portions thereof, including modifications, that you make.

THIS SOFTWARE IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, COPYRIGHT HOLDERS MAKE NO REPRESENTATION OR WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS. COPYRIGHT HOLDERS WILL BEAR NO LIABILITY FOR ANY USE OF THIS SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

CERN ACKNOWLEDGEMENT

This product includes computer software created and made available by CERN. This acknowledgment shall be mentioned in full in any product which includes the CERN computer software included herein or parts thereof.

Oracle WebServer 2.0.2 contains encryption and/or authentication engines from RSA Data Security, Inc. Copyright 1996 RSA Data Security, Inc. All rights reserved.

All other products or company names are used for identification purposes only, and may be trademarks of their respective owners.

Preface

The Oracle WebServer User's Guide is part of the Oracle WebServer documentation set, which contains:

- This book, the *Oracle WebServer User's Guide*
- The Oracle WebServer online documentation
- The *Oracle WebServer Quick Reference*
- The *Oracle WebServer Installation Guide* for your platform.

The online documentation is provided in both `html` and `pdf` formats. To view the `html` files, you can use any standard web browser program that supports tables. Once you have installed the Oracle WebServer, you can follow the link from the product home page to the online documentation.

This Preface discusses this Guide's:

- Organization
- Typographic conventions
- Related documents

How this Guide is Organized

- Chapter 1, "Oracle WebServer Concepts" defines and explains WebServer terms and concepts, and provides background information. This chapter gives an orientation to the product.

- Chapter 2, "Using the Oracle WebServer Manager" provides a brief introduction to the Oracle WebServer Manager and lists the tasks you can perform with it.
- Chapter 3, "Setting Up a Secure Oracle WebServer" describes how to configure your Oracle WebServer to accept secure connections using the Secure Sockets Layer (SSL).
- Chapter 4, "Developing Applications for the Oracle WebServer" describes in detail how to develop applications to run on the Oracle WebServer.
- Chapter 5, "Sample Applications" provides an overview of the sample WebServer back-end applications provided with the documentation.
- Chapter 6, "The PL/SQL Web Toolkit Reference" provides a detailed reference for the PL/SQL procedures you can use to write programs that generate HTML documents dynamically.
- Chapter 7, "Oracle WebServer Messages" provides a lookup reference for errors and messages that the Oracle WebServer can display to the user or log to files.
- Appendix A, "Glossary" provides a lookup reference for terms and concepts used throughout this book.
- Appendix B, "Overview of the Oracle7 Server, SQL, and PL/SQL" summarizes the capabilities of the Oracle7 database server and provides a syntax summary of the SQL and PL/SQL languages.
- Appendix C, "Introduction To HTML" summarizes the capabilities of HTML, gives usage examples, and provides a brief syntax reference.

Conventions Used in This Manual

Feature	Example	Explanation
monospace	`enum`	Identifies code elements.
boldface	**mna.h** **timeout**	Identifies file names and function arguments when used in text.
italics	*file1*	Identifies a place holder in command or function call syntax; replace this place holder with a specific value or string.
ellipses	n,...	Indicates that the preceding item can be repeated any number of times.

Table 1: Conventions

Example Conventions

This Guide shows code in this font:

```
applet.addParam("text", "This is an applet test.");
```

Related Documents

Part No.	Document Title
A44047-1	Oracle WebServer Quick Reference

Table 2: Related Documents

Your Comments Are Welcome

We value and appreciate your comments as an Oracle user and reader of the manuals. As we write, revise, and evaluate our documentation, your opinions are the most important input we receive. At the back of our printed manuals is a Reader's Comment Form, which we encourage you to use to tell us what you like and dislike about this manual or other Oracle manuals. If the form is not available, please use the following address or FAX number.

Oracle WebServer Documentation Manager
Oracle Corporation
500 Oracle Parkway
Redwood City, CA 94065
U.S.A.
FAX: 415-506-7200

Contents

Oracle WebServer Concepts 1-1
Overview 1-2
The Web Listener 1-3
The Secure Sockets Layer 1-7
The Web Server Manager 1-9
The CGI Interface 1-10
The Web Request Broker (WRB) 1-11
PL/SQL Agent 1-13
Java 1-15
LiveHTML 1-16

Using the Oracle WebServer Manager 2-1

Setting Up a Secure Oracle WebServer 3-1

Developing Applications for the Oracle WebServer 4-1
Application Development - an Overview 4-1
Server Extensions 4-2
The PL/SQL Agent 4-4
The Java™ Interpreter 4-15
The LiveHTML Interpreter 4-30

Sample Applications 5-1

The PL/SQL Web Toolkit Reference **6-1**
Installing the Oracle WebServer PL/SQL Web Toolkit 6-3
Procedure and Function Reference 6-5
Head Related Tags 6-10
Body Tags 6-13
Frame Tags 6-29
List Tags 6-31
Character Format Tags 6-37
Physical Format Tags 6-41
Form Tags 6-42
Table Tags 6-51
OWA_UTIL Package 6-55
OWA_PATTERN Package 6-63
OWA_TEXT 6-75
OWA_IMAGE 6-78
OWA_COOKIE 6-79
OWA_INIT 6-81

Oracle WebServer Messages **7-1**
00001 - 00600 Generic Oracle WebServer Configuration Messages 7-1
5000 - 5499 : Oracle Web Agent errors 7-3
5500 - 5599 : Oracle Web Agent Administration errors and messages 7-5
5600 - 5699 : Oracle Web Database Administration errors 7-7
5700 - 5799 : Oracle Web Listener Configuration errors and messages 7-8
5800 - 5899 : Oracle Web Request Broker Administration messages 7-10
5900 - 5999 : Oracle WebServer Registration errors and messages 7-11
7500 - 7599 : Oracle Web Server Proxy errors and messages 7-11
Oracle Java Web Toolkit Messages 7-11

Glossary **A-1**

Overview of the Oracle7 Server, SQL, and PL/SQL **B-1**
Oracle7 Server B-1
SQL B-5
PL/SQL B-12
Basic Structure and Syntax B-14
The DECLARE Section B-14
The EXECUTABLE Section B-21
The EXCEPTION Section B-25
Storing Procedures and Functions in the Database B-28
Database Triggers B-33

Introduction To HTML C-1
What is HTML? C-2
Getting Started C-3
Document Structure C-4
Body Tags C-5
List Tags C-8
Hypertext Linking C-10
Reviewing Changes to Your HTML Document C-15
Adding Style to Your HTML Document C-15
Special HTML Tags C-16
Tables C-17
Forms C-20
Creating Your Own HTML Document C-25
More Information about HTML C-26



CHAPTER

1

Oracle WebServer Concepts

This section offers a conceptual overview of the Oracle WebServer. It is intended to give you a context for the more specific task-oriented information that follows. The subsequent sections of this book are specifically targeted at those who manage Oracle WebServer sites (WebServer Administrators) and those who write application programs for them (application developers). This section covers:

- *Overview*.
- *The Web Listener*. This is the portion of the WebServer that interfaces to the local network or the World Wide Web.
- *The Secure Sockets Layer*. Introduction to SSL public key encryption.
- *The Web Server Manager*. The set of Web pages you can use to perform most WebServer administration.
- *The CGI Interface*. The standard Web mechanism for executing applications on a Web server.
- *The Web Request Broker (WRB)*. The core of the WebServer. An asynchronous request broker with an open API that Oracle WebServer uses to execute applications on the server.

- The *PL/SQL Agent*. The program the Oracle WebServer uses to execute procedures written in PL/SQL, Oracle's application development language, on the Oracle7 Server.
- *Java*. A new language for developing distributed network applications. Oracle WebServer enables you to execute Java either on the WebServer itself or on the client's browser.
- *LiveHTML*. A way to embed dynamic content in Web pages. This content can be either other Web pages or the output of scripts run by the Operating System. LiveHTML is an Oracle extension of the NCSA standard Server Side Includes functionality.

Overview

The Oracle WebServer is a HyperText Transfer Protocol (HTTP) Internet Server with unprecedented database integration and a powerful development environment. When the WebServer receives a Uniform Resource Locator (URL) from a browser located either on the World Wide Web or on a local network using the Web's protocol (HTTP), it draws on information from the database and the operating system's (OS) file system as necessary to respond to the request. The file system can be used for static (hardcoded) Web pages, or for CGI scripts that do not access the database, and the database is used for Web pages that are generated at runtime using "live" data. Although you can run the Oracle WebServer without Oracle7, one of the great advantages of the product is its tight integration with the Oracle7 Server, the leading database product in the world.

The Web Listener is the component that receives a URL from a Web browser and sends back the appropriate output. When the Web Listener receives a URL, it determines whether the request requires the use of a Service to be accessed through the Web Request Broker (WRB), a program to be accessed through the CGI interface, or whether access to the file system of the machine on which the Listener resides is sufficient. If WRB access is required, the Listener passes the request to *WRB Dispatcher* for processing; then it returns to the task of listening for more incoming HTTP requests.

The WRB Dispatcher handles requests with the aid of a pool of processes called *WRB Executable Engines (WRBXs)*. Each WRBX interfaces to a back-end application using the WRB API. These applications are called *WRB cartridges*. The WRB API is designed so third parties can add their own cartridges. The combination of a cartridge and the WRB API is known as a *WRB Service*. Currently, Oracle WebServer support three kinds of WRB Services:

- PL/SQL cartridges. These execute stored PL/SQL procedures to generate HTML dynamically using Oracle data.
- Java™ cartridges. These execute Java code on the Server.
- LiveHTML. These embed Web pages within one another, and store in Web pages the output of scripts executed by the Operating System.

The Web Listener

The Oracle Web Listener is a HyperText Transfer Protocol (HTTP) engine that responds to requests for hypermedia documents from web browsers (clients). This section summarizes the Web Listener's capabilities.

Network Communication

Every Oracle Web Listener process accepts connections from web browsers (clients) on one or more IP address/port combinations using HTTP to encode requests for hypertext documents, and the Transmission Control Protocol/ Internet Protocol (TCP/IP) as the underlying connection protocol. Several Web Listener processes may run on single host computer at the same time (see the *Oracle Web Listener Administration Form* for more information).

Ports

A *port* is a number that TCP uses to identify a communication channel associated with a specific program. For example, the login program usually accepts login requests on port 49, and the Domain Name Service (DNS) usually accepts name lookup requests on port 53. HTTP engines, such as the Oracle Web Listener, usually accept ordinary connections on port 80 and secure connections on port 443.

Programs must execute with root permissions to access port numbers 1 through 1023, whereas any program can access port numbers 1024 through 65535.

IP Addresses

An Internet Protocol (IP) address is a unique number that identifies exactly one computer on the Internet, although each computer can be represented on the Internet by several addresses. A computer can use different addresses for different Internet functions. For example, a single computer might have one IP address on which it acts as an email routing server, and another address on which it acts as a DNS server.

DNS Host and Domain Names

A *fully qualified host name* is a unique character string that identifies exactly one computer on the Internet, although each computer can be represented on the Internet by several host names. The Domain Name Service (DNS) maintains a distributed database that maps host names to IP addresses. *DNS servers* provide Internet computers with lookup services for this database. `hal.us.oracle.com` is an example of a fully qualified host name.

A *domain* is a named DNS host name space, such as `us.oracle.com`. All host names within a domain must be unique. For example, there must be only one host named `hal` in the `us.oracle.com` domain.

Secure Connections

The Oracle Web Listener can accept secure connections using the Secure Sockets Layer (SSL), an emerging standard for secure data transmission. (The combination of IP address and port number is called a *socket*.) See *The Secure Sockets Layer* for more information on the Oracle WebServer's implementation of SSL.

Clients use a secure version of HTTP called HTTPS to establish secure connections with SSL. To establish a secure connection with a Web Listener process, a request URL must use "https" instead of "http" must specify a port on which the Web Listener has enabled SSL, for example:

```
https://www.blob.com:443/
```

Proxies

You can configure an Oracle Web Listener process to act as a *proxy server*, an HTTP engine running on a firewall machine that allows clients inside the firewall to access web sites outside the firewall. The Oracle Web Listener implements the proxy behavior defined by Luotonen and Altis in *World-Wide Web Proxies* (http://www.w3.org:80/hypertext/WWW/Proxies/) under the auspices of the World-Wide Web Consortium (W3C).

File Handling

The Oracle Web Listener uses a *virtual file system* to keep track of the files it makes available to clients. The virtual file system maps the pathnames used in request URLs (Uniform Resource Locators) to the file system maintained by the host machine's operating system.

File Memory Mapping

The Web Listener uses the host operating system's memory mapping capability when opening a file requested by a client so that the Web Listener's virtual address space refers directly to the file's contents. This speeds access, and makes it possible for several clients to share the same copy of the file, which conserves the Web Listener's memory resources.

File Caching

Ordinarily, when the Web Listener opens a requested file, the file remains open and mapped into memory until all clients using the file are finished with it, at which point the Web Listener closes the file and frees the memory mapping associated with it. The Web Listener allows you to specify files to be *cached*. Cached files are opened when a client requests them, and remain open for the life of the Web Listener process. This optimizes access times for files, such as your home page, that are requested often.

File Protection

The Oracle Web Listener allows you to make secure specific virtual files or directories by assigning authentication and/or restriction schemes to protect them.

Authentication Schemes

When a file or directory is protected by an *authentication scheme*, a client requesting access to it must provide a user name and password. Authentication schemes allow you to define named groups of user name/password combinations, and named *realms* that are groups of these groups. You can then assign user, group, and realm names to virtual files and directories, requiring any client requesting access to input one of the specified user name/password combinations.

The Web Listener supports two authentication schemes: *basic authentication* and *digest authentication*. Both schemes are identical, except that digest authentication transmits passwords from client to server in an encrypted form called a *digest*, whereas basic authentication sends unencrypted passwords, making it considerably less secure.

Some older web browsers don't support digest authentication, but for files or directories that require authentication, you should use digest authentication whenever possible.

Restriction Schemes

When a file or directory is protected by a *restriction scheme*, only a client accessing the Web Listener from a trusted group of host machines may access it. The two restriction schemes that the Web Listener supports are *IP-based restriction* and *Domain-based restriction*. IP-based restriction allows you to define groups of trusted hosts identified by IP address, whereas Domain-based restriction allows you to define groups of trusted hosts identified by DNS host or domain name.

File Format Negotiation

The Oracle Web Listener can maintain several versions of a document in different formats, and provide it to clients in the formats they prefer.

Language Formats

For example, if a client requests a document in French, the Web Listener checks to see if it has a French version of the document, and if so, returns that version to the client. Otherwise, the Web Listener returns the version of the document in its own default language (usually English).

MIME Formats

Similarly, a client can request a document of a particular Multipurpose Internet Mail Extensions (MIME) type. If the Web Listener can find a version of the requested file in the requested MIME format, it returns that version to the client. Otherwise, it returns the version of the file that has the smallest size.

Encodings

The Web Listener can also maintain versions of a document that have been encoded by programs such as a compression utilities. For example, if a client can uncompress a document that has been compressed by the gzip program, the Web Listener can return to the client the compressed version of document instead of the uncompressed document, saving transfer time.

Filename Extensions

The Web Listener uses filename extensions to identify a file's format, which can represent language, MIME type, or encoding. For ease of maintenance, it's best to advertise a file to clients only by its base name, allowing clients and server to negotiate formats transparently.

Dynamic Document Generation

Using the Oracle Web Listener, your web site can respond to client requests by generating HTML documents dynamically. This allows you to customize your WebServer's responses.

Like most HTTP engines, the Oracle Web Listener allows clients to use the Common Gateway Interface (CGI) to run programs on the server machine to perform special processing and return data to the client.

The Web Request Broker

Unlike other HTTP engines, the Oracle Web Listener provides an interface called the Oracle Web Request Broker (WRB), which allows clients to run programs on the server machine and return data much more efficiently than CGI allows. To do this, the Web Listener passes requests intended for these programs to the WRB Dispatcher, which maintains a pool of processes to which it can assign the requests. See *The Web Request Broker (WRB)* for more information.

The Secure Sockets Layer

The Secure Sockets Layer (SSL) is an emerging standard for secure data transmission over the Internet.

One problem with communicating sensitive information over the Internet is that almost every connection between two computers over a network involves many intermediate steps—a chain of computers that successively receive and forward the information until it reaches its destination. This process, called *routing*, is fundamental to all Internet communication, and any computer in the routing chain has complete access to all the data it receives.

This makes it easy for the unscrupulous to intercept your private conversations, steal your credit card numbers, or illegally obtain confidential or proprietary information.

The Oracle WebServer's implementation of SSL addresses this problem by scrambling data sent from the server to clients (web browser programs) in such a way that the clients can unscramble the information when they receive it. This way, any intermediate computers involved in routing the information see only gibberish that they can't decipher.

This kind of security has three aspects:

- Encryption—a mechanism for scrambling and unscrambling data.

ORCL000800

- Authentication—a mechanism by which the one party proves its identity to another party.
- Data integrity—a mechanism for verifying that all of the data transmitted, and *only* the data transmitted, is received correctly.

Encryption

A traditional encryption system, called a *secret-key* system, uses a single large number called a *key* both to scramble (encrypt) and unscramble (decrypt) messages. Secret-key encryption systems are very fast, but they rely on one party communicating the secret key to another party, often by way of a third party such as a courier, before the two parties can exchange encrypted messages. This makes keys vulnerable to theft or tampering while in transit.

Public-Key Encryption

To avoid this problem, SSL uses a form of encryption called *public-key encryption* to encrypt and decrypt transmitted data. Unlike secret-key encryption systems, a public-key system uses pairs of keys (*key pairs*). One key, called the *public key*, is used to encrypt messages, while the other, called the *private key*, is used to decrypt messages. The two keys are large numbers that are related mathematically in such a way that it takes a very long time to calculate the private key from the public key.

If you want to receive encrypted messages using public-key encryption, you must first run a program that generates a key pair. You must then publish the public key in a public database or directory, and store the private key in a secure location on your computer. *This is critical.* The effectiveness of public-key encryption depends entirely on the secrecy of the private key.

Anyone who wants to send you an encrypted message must look up your public key in a directory, use it encrypt the message, and send you the encrypted message. Only your private key can decrypt the message, so if you have kept your private key secret, no one else can read the message.

Because public key encryption is much slower than secret-key encryption, SSL uses it only when the client first connects to the WebServer to exchange a secret key called a *session key*, which both client and server use to encrypt and decrypt transmitted data.

Authentication

Another application of encryption is authentication. Authentication using public-key encryption involves using a *digital signature*, an electronic proof of identity analogous to a handwritten signature.

If you want to "sign" an electronic document in a verifiable and legally binding way, you must first possess a key pair. You must then run a program that generates a digital signature using the private key and the document itself. You can then attach the digital signature to the document and send it. Anyone who receives this document, together with its digital signature, can then use the your public key to verify your identity, and to verify that the document has not been tampered with.

Certificates and Certifying Authorities

When clients connect to your web site for a transactions that require them to transmit sensitive information, they must be assured that they haven't connected to an impostor pretending to be you. Clients therefore require your WebServer to authenticate itself before such transactions can proceed.

To authenticate itself, your WebServer must present the client with the proper credentials, called a *certificate*.

When you set up a secure WebServer, you must obtain a certificate from a trusted third-party company called a *certifying authority* (CA).

When you contact a certifying authority to request a certificate, you must provide them with certain legal information about your organization, which they can use to certify that your organization is legitimate and should be certified (see *Setting Up a Secure Oracle WebServer* in the online documentation).

The Web Server Manager

To help you manage your Web site, the Oracle WebServer provides a set of Web pages that you can use to perform most common administration tasks. These pages are simply an easy way to edit the configuration files that the WebServer uses, so you can always use another tool to edit those files directly, although Oracle discourages this.

The WebServer Manager can be used from any Web browser. It has the following sections:

- *The Listener Pages*

- *The WRB Pages*
- *The PL/SQL Agent Pages*
- *The Oracle7 Server Pages*

The Listener Pages

The largest group of pages deal with administering the *Web Listener*. For the most part, however, these pages are just forms where you fill in the values for various configuration parameters. The exception is the security pages, where you specify which sorts of security schemes will be used, including the following:

- If you select Restriction, the IP addresses or domain names that are granted various levels of access.
- If you select Authentication, the user names and passwords that are granted various levels of access.

The WRB Pages

These pages let you specify the actual and virtual directories for WRB cartridges, as well as the number of WRBXs to assign to each.

The PL/SQL Agent Pages

Administering the *PL/SQL Agent* means administering the *Database Connection Descriptors (DCDs)* that it uses to establish its identity when communicating with the database.

The Oracle7 Server Pages

The *Oracle7 Server* is an extremely sophisticated product, and it is properly administered using Oracle Server Manager or directly through the SQL language. The pages provided here enable you to do a few basic things, such as start up and shut down database instances, and browse database contents.

The CGI Interface

The Common Gateway Interface (CGI) is the standard technique used by an HTTP server to execute a program that generates HTML output. Using CGI, you can run *PL/SQL*, and thus interface to the Oracle7 Server, from Internet servers

that do not support the *WRB*. This technique provides dynamic content rather than static content from files on disk. Oracle WebServer is fully compliant with CGI version 1.1.

When the Web Listener recognizes an incoming URL as a request to execute a CGI application, it spawns a separate process to perform the operation (*The Web Request Broker (WRB)* circumvents this need to spawn a new process for each request, thereby improving performance). The Web Listener passes the URL to the process it spawned and maintains communication with it through standard input and output. Therefore, the CGI process can get the input it needs from the URL itself and/or the standard input. It sends its output back to the Listener through the standard output, and the Listener transmits it in turn to the client's Web browser.

The fact that CGI applications spawn a new process each time they are used is costly in terms of performance. For this among other reasons, Oracle recommends that you use the WRB instead.

The Web Request Broker (WRB)

The WRB (Web Request Broker) is an asynchronous request handler with an API (Application Program Interface) that enables it to interface dynamically and seamlessly to various back-end technologies called "WRB Services".

Whenever the Web Listener receives a URL that calls for the WRB, it passes execution of the request to the *WRB Dispatcher* or simply *Dispatcher*. The Dispatcher maintains communication with a pool of processes called *WRB Executable Engines (WRBXs)*.The Dispatcher finds a free WRBX that is configured to run the desired WRB Service and passes execution to it. The result is that the Listener can receive and validate URLs coming in, while each request is handed off to a process that executes it in the background. Each WRBX uses the WRB API to interface to a WRB *cartridge*. WRB cartridges can be of the following types:

- *The PL/SQL Agent*. This cartridge executes PL/SQL commands stored in the database. It is better optimized for database access than the Java cartridge, but doesn't have all of Java's functionality.
- *The Java™ Interpreter*. This cartridge lets you execute Java on the server to generate dynamic Web pages. You can also execute PL/SQL from within Java using this cartridge.
- *The LiveHTML Interpreter*. This cartridge is Oracle's implementation and extension of the industry-standard Server Side Includes functionality.

LiveHTML enables you to include in your Web pages the output of any program that your Operating System can execute.

WRB Services

The combination of a cartridge and its WRB API constitute the *WRB Service*. WRBXs are instances of WRB Services, so that there are three WRB Services corresponding to the three WRB cartridges, while the WRBXs are created and destroyed according to the workload.

The Dispatcher determines which WRB Service to send a given request to on the basis of the path and the file extension (MIME type) specified in the URL. You configure the WRB in the WebServer Manager to determine which combinations of path and extension correspond to which Services.

The WRB API is designed so that third parties can write their own extensions to the WebServer. The three cartridges listed above come with the WebServer, but you can obtain additional cartridges from third parties or write your own.

The WRB Services can also augment the access control specified by the Listener with Service-specific control. The Services included with the WebServer do not do this, but the WRB API provides a callback routine that enables you to put this functionality in cartridges that you write. This routine enables your cartridge to do any of the following:

- Prompt for a username and password, using either a Basic or a Digest scheme. The difference between Basic and Digest is that a Digest scheme encrypts the password (for more information on these terms, see the discussion of the Listener earlier in this section).
- Specify by name some restriction scheme to be applied to this request. Such restriction schemes are defined in the Listener and limit access by IP address or Domain Name.

For more information on how WRB cartridges work and how to write your own, refer to the WRB specification that accompanies the WebServer online documentation.

WRB Executable Engines

The Dispatcher creates and maintains the WRBXs. The data passed to a WRBX when a request comes in consists of:

- The URL triggering the request. This can include a file extension indicating the desired MIME type of the result.

- The desired language of the result.
- The desired character set of the result (how the language is encoded, for example, ISO or Unicode).
- CGI environment variables, which allow WRB cartridges to be run by applications written for *CGI*.
- If the request involves the use of the PL/SQL Agent, which of the *Database Connection Descriptors (DCDs)* to use.

The Dispatcher continually adjusts the load by controlling how many WRBXs are running at a given time, subject to certain parameters. These parameters are set by the WebServer Administrator, who sets the maximum and minimum number of WRBXs running for each WRB Service. The Dispatcher creates new WRBXs as needed and connects them to the appropriate WRB Services. The Dispatcher also keeps track of which WRBXs are executing requests and which are free. The WRBXs are single-threaded processes that communicate with the WRB Dispatcher using the dataflow mechanism appropriate to the Operating System (such as a pipe in Unix).

PL/SQL Agent

The PL/SQL Agent can be invoked through either the WRB or CGI, as determined by the directory and MIME type mappings set by the WebServer Administrator. If it is through CGI, the script name specified in the URL must be "owa". If it is through the WRB, the directory and MIME type mappings are sufficient to specify the PL/SQL Agent. The PL/SQL Agent executes application code written in PL/SQL and returns the output in HTML form for the Web Listener to output as a Web page.

PL/SQL procedures are stored in the database. The PL/SQL Agent invokes them by issuing commands to the database, which then performs the actual execution and sends the output and status messages back to the PL/SQL Agent.

Since PL/SQL is actually executed in the database, the PL/SQL Agent, whether executed through the WRB or CGI, must connect to the Oracle7 Server to run. If the PL/SQL Agent uses the WRB, the WRBX connects to the database in the following two stages:

- It establishes the connection. The WRBX does this as soon as it is created.
- It logs on to the database. This is a separate operation, and the WRBX does not do this until a request comes in, so as to give each request a separate

database session. Once the request is handled, the WRBX logs off of the database, but keeps the connection intact.

Since the first stage actually takes most of the time involved in establishing a database session, this technique speeds execution considerably, yet the effect is as though each request connected to the database independently. If run through CGI, the PL/SQL Agent must go through the entire procedure for each request. The username, schema, and password that the PL/SQL Agent uses to connect is specified by the URL through the use of a DCD, as explained below.

Specifying the Database Connection

A URL that invokes the PL/SQL Agent must specify a *DCD* (Database Connection Descriptor). This is an OS file maintained by the WebServer that provides the username, password, database, and other information to be used to establish the database connection. The DCD determines both the database access privileges the PL/SQL Agent has when executing this request and the schema (portion of the database) that it accesses. The filename of the DCD is given in the URL as a file within the directory configured for the PL/SQL Agent. The file need not actually reside in that directory; its association with the PL/SQL Agent is set in the WebServer Manager PL/SQL Agent configuration page. If the PL/SQL Agent is invoked through CGI, the DCD precedes the script name "owa", with the two separated by a slash (/).

Oracle WebServer also provides you with Java classes that can invoke PL/SQL. For more information, see Java or the Java Interpreter.

The PL/SQL WebToolkit

To make it easier for you to develop Web applications using Oracle data, Oracle WebServer provides you a group of *PL/SQL* packages that you can use to easily generate Web pages from data stored in an Oracle database. These packages are called the PL/SQL Web Toolkit. The intent is for you to create PL/SQL procedures that access and process the Oracle data you wish to place on the Web. From within these procedures, you call the PL/SQL Web Toolkit procedures you need to create the HTML you want. You store the procedures you write in the database, just as other PL/SQL packages, including the toolkit, are stored. You also design your Web pages, including the dynamically-generated ones, to produce URLs that call the PL/SQL procedures you want in response to specified user actions. Having your code executed within the database brings many performance, security, and portability benefits. For more information on the PL/SQL Web Toolkit, see *The PL/SQL Web Toolkit Reference*.

Java

Java is an object-oriented language for creating distributed applications on the Internet or other networks. Modules of Java code known as "applets" can be downloaded from the Internet or a local network in real time and locally executed. Java applets themselves can call and execute other applets, so that a Java application as executed on the user's machine can be constructed "on the fly" from a repertoire of standard parts that reside on the net. You might call this the "building block" approach to programming. Among the important features of Java are the following:

- It is extremely portable. Java code is compiled to a form known as "bytecode". This is a sort of generalized computer code that is not executable by any particular machine, but is recognized by the "Java Virtual Machine". It is as bytecode that Java applets transverse the net. The Java Virtual Machine (VM) resides on the computer where the applet is to be executed and converts the bytecode to the native code for that machine. Currently, Java VMs exist or are planned for all current versions of Windows, Solaris, MacOS, OS/2, Linux, Amiga, and other platforms.
- It is fully object-oriented. Languages like C++ that add object-oriented features to non-object-oriented languages must make compromises. Java applets do not allow violation of object-oriented principles such as "encapsulation".
- It uses a C-like syntax. This makes the language easier for C programmers to learn.
- It is multi-threaded, in effect executing several chains of control flow concurrently. The Java language itself provides tools for managing the threads, rather than relying exclusively on the OS.
- It prohibits direct memory manipulation. In Java, there are no pointers and no direct memory allocation. This eliminates a rich source of C's functionality, and an even richer source of its bugs.
- You can embed calls to Java applets in Web pages, and the applet will be executed by the browser, provided it is Java-enabled. Most major browsers plan to support Java.

Note: Though powerful, Java is a young technology. Oracle WebServer supports it because of its rich features and wide acceptance. However, you should be aware that it may be somewhat less stable than more mature technologies.

Client vs. Server Side Java

Oracle WebServer supports the use of Java either on the client, which is to say any Java-enabled Web browser, or on the server. Code to be executed on the client is for the most part extracted and manipulated like other data. The best way to handle such code is to store it in the OS file system and extract it in real time.

You can also execute Java as a *WRB* cartridge on the WebServer itself. You might want to do this, for example, to perform graphical manipulation for which PL/SQL is ill-suited. For example, you can combine several graphics from the database into a single image. Each region of the image would be a separate button that the user can click, and each button clicked would produce a different effect. In HTML, this is called an "image map". Using Java on the server, you could generate such image maps dynamically, with the components of the image being based on the results of a database query.

To execute Java on the server, you use the WRB API to interface directly to the *Java Interpreter* residing in the WebServer. This interpreter finds and executes the Java code and returns the results, through the WRB interface, to the Web Listener.

To make it easier for you to develop Java applications, Oracle WebServer provides the Java Web Toolkit, a group of Java packages containing classes to aid in database access and dynamic HTML generation.

Using PL/SQL Within Java

Since *PL/SQL* code is actually part of the database, you can call it from within Java, which enables you to create applications that combine the strengths of both languages. Because PL/SQL execution takes places in the database, doing this does not hinder the portability of the application. A PL/SQL application can execute without modification on any platform where the *Oracle7 Server* runs, just as a Java applet can execute without modification on any platform that has a Java Virtual Machine (VM).

LiveHTML

LiveHTML is Oracle's implementation and extension of the standard Server Side Includes functionality defined by the NCSA. *The LiveHTML Interpreter* enables you to include dynamic content in otherwise static Web pages. At the point in

your Web page where you want to interject dynamic content, you place a tag that points to one of the following:

- A static Web page.
- Another LiveHTML Web page.
- A script that is executed on the server and outputs HTML. This script may but need not conform to the *CGI* standard.
- A system variable, for example: FMODDATE.

You can use a variable to determine at runtime the Web page, variable, or script to which the tag points. This enables you to have a Web page that selects dynamically from among any number of static Web pages or scripts, based, for example, on values a user provides in an HTML form. The result is a dynamic Web page built of static Web page components, variables, and HTML output from scripts.

A Web page that is to use LiveHTML must be parsed by the WebServer. For this reason, it differs slightly from ordinary Web pages written in HTML, which the WebServer simply delivers to the browser. To have the LiveHTML tags executed on the server, you must use the WebServer Manager to specify that a given Web Listener is to parse files for LiveHTML. You have the option of having the Listener parse all files or just those with certain extensions.

Enabling users to execute scripts on the server can create security and other risks. For this reason, you can specify that a specific Listener allows only "crippled" includes. This means that LiveHTML parsed by that Listener will be able only to call static HTML, environment variables, or other server parsable files, not executable scripts.

You frequently use LiveHTML when you have standard components, such as menus, that you want on many pages. If desired, you can use LiveHTML to run the PL/SQL Agent under *CGI*, and thereby incorporate dynamic Oracle data in hardcoded Web pages.



CHAPTER

2 Using the Oracle WebServer Manager

The Oracle WebServer Manager is a collection of HTML forms you can use to configure your Oracle WebServer. The following links go to WebServer Manager forms that allow you to:

- Start and stop Oracle7 databases.
- Configure Oracle Web Listener processes.
- Configure the PL/SQL Agent.
- Configure the Web Request Broker.

Starting and Stopping Oracle7 Databases

The *Oracle WebServer Database Administration Form* provides a convenient way to start and stop Oracle7 databases that your WebServer uses. To perform other database administration tasks, see the *Oracle7 Server Administrator's Guide.*

Configuring Oracle Web Listener Processes

The Oracle Web Listener is the Oracle WebServer's HTTP engine, which handles connections from clients (web browsers). You can run several Web Listener processes at once on your WebServer host computer. Using the *Oracle Web Listener Administration Form*, you can create and configure a new Web Listener

process, or choose an existing Web Listener process from a list and either modify its configuration, or delete it.

Configuring the PL/SQL Agent

The PL/SQL Agent allows the Oracle Web Listener to access an Oracle7 database in response to an HTTP request, and return the results to the requestor in HTML form. There are two versions of the PL/SQL Agent:

- The PL/SQL Agent WRB Service—uses the Web Request Broker (WRB) interface to communicate with the Web Listener.
- The PL/SQL Agent CGI Program—uses the Common Gateway Interface (CGI) to communicate with the Web Listener.

The PL/SQL Agent WRB Service is the preferred version because it takes advantage of the high performance of the Web Request Broker interface. The PL/SQL Agent CGI Program is provided only for backward compatibility to version 1.0 of the Oracle WebServer.

Configuring the Web Request Broker

The Web Request Broker provides a programming interface for writing WebServer back-end applications. You can use the Web Request Broker Administration Form to perform general Web Request Broker configuration, and to configure specific WRB cartridges.

Oracle WebServer Manager Tasks

The following list is a summary of the configuration tasks you can perform using the Oracle WebServer Manager. Following one of these task category links will take you to a form where you can perform the tasks in that category:

Oracle7 Database Administration

- Starting up an Oracle7 database
- Shutting down an Oracle7 database

Oracle Web Listener Administration

- Creating a new Web Listener
- Modifying a Web Listener configuration
- Configuring Web Listener network parameters
- Configuring Web Listener logging parameters

- Configuring Web Listener user and group parameters
- Configuring Web Listener virtual directory mappings
- Configuring Web Listener file caching
- Defining Web Listener file extensions for language formats
- Defining Web Listener file extensions for MIME types
- Defining Web Listener file extensions for encoding formats
- Configuring Web Listener security
- Controlling access to virtual files and directories
- Configuring port security (SSL)
- Starting a Web Listener process
- Stopping a Web Listener process
- Deleting a Web Listener
- Enabling Web Listener proxy behavior
- Enabling LiveHTML parsing for a Web Listener

PL/SQL Agent Administration

- Creating a new Database Connection Descriptor (DCD)
- Creating a default DCD
- Creating a DCD for database administration
- Creating a new DCD from an existing configuration
- Modifying a DCD
- Deleting a DCD

Web Request Broker Administration

- Configuring the Web Request Broker
- Configuring a new WRB cartridge
- Modifying a WRB cartridge configuration
- Deleting a WRB cartridge





CHAPTER

3 Setting Up a Secure Oracle WebServer

You can up an Oracle WebServer process to accept secure connections on a particular TCP/IP port by configuring it to use the Secure Sockets Layer (SSL) on that port. SSL is an emerging standard for encrypted data transmission—see *The Secure Sockets Layer* for an introduction to the terms and concepts involved in this kind of security.

To set up your Oracle WebServer to use SSL, you must do the following:

1. Generate a certificate request.
2. Send the request to VeriSign, Inc., a Certifying Authority (CA).
3. Physically secure and prepare your WebServer host machine.
4. Install the certificate granted you by the CA.
5. Activate SSL on at least one WebServer port.

Generating a Certificate Request

To generate a certificate request, run the interactive utility `genreq` and enter the information for which it prompts you. When the prompt specifies a default value, you can just press return to enter that value, or enter a different value if you prefer. For an example of how to use genreq, see the sample `genreq` session that accompanies this document. To run `genreq`, do the following:

1. Type `genreq` to start the utility.
2. Type `G` to begin creating a certificate request.
3. When prompted, type a password, used in generating the private key. Just choose a random string of characters—you need only remember this string long enough to repeat it in the next step.
4. Retype the password for confirmation.
5. Choose the public exponent you want to use in generating the key pair. The only two recognized exponents are 3 and 65537, commonly called Fermat 4 or F4.
6. Enter the size in bits of the modulus you want to use in generating the key pair. For the version of genreq sold in the United States of America, the size may be from 1 to 1024. The default size is 768 bits and the maximum is 1024 bits. A modulus size between these two values is recommended.

   For versions of `genreq` sold outside the USA, the maximum (and default) modulus size is 512 bits.
7. Choose one of three methods for generating a random seed to use in generating the key pair:
   - F—`genreq` prompts you to enter the full pathname of a file in your local file system. This can be any file that is at least 256 bytes in size, does not contain any secret information, and has contents that cannot easily be guessed.
   - K—`genreq` prompts you to enter random keystrokes. `genreq` uses the variation in time between keystrokes to generate the seed. Don't use the keyboard's autorepeat capability, and don't wait longer than two seconds between keystrokes. `genreq` prompts you when you have typed enough keystrokes. You must delete any unused characters typed after this prompt.
   - B—genreq prompts you to enter both a file name and random keystrokes. This option is recommended.
8. Enter the name of a file in which to store your WebServer's distinguished name. You can choose the default, or enter any filename with a `.der` extension. `genreq` creates this file in the current directory, though you may later move it to any convenient location.
9. Enter the name of a file in which to store your WebServer's private key. You can choose the default, or enter any filename with a `.der` extension. `genreq` creates this file in the current directory, though you may later move it to any convenient location.
10. Enter the name of a file in which to store the certificate request. You can choose the default, or enter any filename with a `.pkc` extension.

11. Enter the requested identification information for you organization:

- Common Name—the fully qualified host name of your organization's Internet point of presence as defined by the Domain Name Service (DNS), for example, `www.oracle.com`.
- Email Address—the email address where the CA can contact you.
- Organization—the official, legal name of your company or organization. Most CAs require you to verify this name by providing official documents, such as a business license.
- Organizational Unit—(optional) the name of the group, division, or other unit of your organization responsible for your Internet presence, or an informal or shortened name for your organization.
- Locality—(optional) the city, principality, or country where your organization is located.
- State or Province—the full name of the state or province where your organization is located. VeriSign does not accept abbreviations.
- Country—the two-character ISO-format abbreviation for the country where your organization is located. The country code for the United States is "US."
- Email Address—the email address where VeriSign can contact you.
- The name and version number of the application for which you are getting the certificate (you should accept the default value).

  When you have entered all the requested information, `genreq` responds with `Thank you`, and processes the data you have entered. When it is finished, it outputs `done` and returns you to the main menu.

12. Type `Q` to quit the program.

Requesting a Certificate

To request a certificate, email the request generated by `genreq` to VeriSign, Inc. a subsidiary of RSA Data Security, Inc.

The certification process can take time, from a few days to several weeks. The more organized and complete your paperwork, the better your chances are for quick certification.

Preparing Your WebServer Host Machine

For your WebServer to be secure as advertised to clients, you must make sure that no unauthorized person has access to your WebServer's host machine. Here are some suggestions:

- Place the machine in a locked server room.
- Limit distribution of keys or combinations to the server room to a few trusted individuals.
- Set up a secure area of the machine's file system that can be accessed only by the root user. This is where you will store your private key and your certificate when you receive it.
- Set a secure root password on the machine, using at least six characters and mixing numbers, legal punctuation marks, and mixed-case letters. Try not to use a character string that is a proper name or a word in any language. Change passwords frequently, and never write a password on paper.
- Strictly limit the programs that are installed and allowed to run on the machine.
- Limit TCP/IP connections to the machine to port 443, the default port for secure connections. Disable all other ports.

To learn how to implement these suggestions on your WebServer machine, see your machine's operating system reference manual.

Installing Your Certificate

When you receive your certificate from VeriSign, you must use the Oracle Web Server Manager to install it.

1. Use your email reader to save the message from the CA containing the certificate to a file with a `.der` extension, such as `cert.der`.
2. Use a text editor to remove the header information before the BEGIN CERTIFICATE line and the footer information after the END CERTIFICATE line. *Do not* delete the BEGIN CERTIFICATE and END CERTIFICATE lines themselves.
3. Using your web browser, go to the *Oracle WebServer Administration* page.
4. Follow the link to the *Oracle Web Listener Administration Form.*
5. Select a Web Listener process from the list and click its Configure button to go to the *Oracle Web Listener Advanced Configuration Form.*
6. Follow the link to the *Oracle Web Listener Configure Security Form.*
7. Go to the Secure Sockets Layer section of the form and follow the instructions.
8. Go back to the Oracle Web Listener Advanced Configuration Form.
9. Go to the Addresses and Ports section.

10. Set the Security pull-down menu to SSL in the entry for at least one port.

    Note: For increased security, it is best to activate SSL only on port 443.

11. Go back to the *Oracle Web Listener Administration Form* and repeat the process for all Web Listener processes that you want to make secure.





CHAPTER

4

Developing Applications for the Oracle WebServer

This section covers the development of applications for the WebServer. It has the following sections:

- *Application Development - an Overview*
- *Server Extensions*
- *The PL/SQL Agent*
- *The Java™ Interpreter*
- *The LiveHTML Interpreter*

Application Development - an Overview

Applications developed for the Oracle WebServer have two general types of components:

- Web pages, whether statically coded or generated at runtime from PL/SQL or Java. These constitute both the user interface and the final product. Using LiveHTML, these can reference applications executed on the server when the page is accessed.

- Code executed on the server. There are several kinds of code that can be so executed, but all of them eventually produce HTML that is sent to the browser as Web pages.

The Web pages you create will themselves generate URLs that determine what other Web pages are retrieved and what code is executed on the server. Therefore, you must write your Web pages to generate URLs that contain information that the server needs to find the desired Web pages or execute the desired code. These are the general ways your Web pages can produce such URLs:

- They can be explicitly given as part of hypertext anchors. In this case, the URL is transmitted whenever the user triggers the hypertext link.
- They can come from a user typing them directly into a search dialog box.
- They can come from HTML forms.

From the viewpoint of the client, all of these techniques specify the retrieval of a Web page. What determines whether a Web page is simply fetched from the file system or generated by some code executed on the server is the location specified in the URL. The WebServer understands that some locations contain files and others reference WRB Services or CGI programs. It is the configuration of the WebServer that determines which is which. You could say that the server extensions s appear to the client as "virtual files."

Server Extensions

There are three basic techniques for executing applications on the server to dynamically generate Web pages:

- *The PL/SQL Agent*
- *The Java™ Interpreter*
- *The LiveHTML Interpreter*

The first two are languages that the WebServer can cause to be executed. The last is an HTML technique for passing execution to another process at a certain point in the scanning of a Web page, and embedding the output of that process in the calling page. LiveHTML can also execute programs, but its programs are written in some language, such as Perl for Unix or Visual Basic for Windows NT, that is executed by the OS itself rather than the WebServer.

Generally speaking, you can combine these techniques. For example, the Java Interpreter can execute PL/SQL in the database.

The PL/SQL Agent and the Java Interpreter can use either of two interfaces: WRB or CGI. Which interface is used is determined by the mapping of the "file" the URL requests; this mapping is part of the WebServer configuration.

How URLs Specify Applications

Once the Listener has determined that a URL specifies that an application is to be executed on the server, it interprets the URL to extract path information and arguments to be passed to the application on start-up.

URLs that specify applications are split into three different parts:

- The virtual path
- The extra path information
- The query string

The syntax is as follows:

```
virtual_path extra_path_information?query_string
```

Here are the explanations of the syntax components:

virtual path

This is similar to a path you would use to access a regular document or image. That is, it is a pathname that identifies the application that you want executed.

extra path information

This is optional additional information embedded in the URL after the pathname. This consists primarily of various environment variables that you can use to pass information to the application.

query string

This is another optional part of the URL. This is used to directly supply parameter values (as opposed to environment variable values) to the requested application. The parameters are specified in the following form:

```
<parameter name>=<value>
```

Example of a URL Invoking an Application

This section provides an example of a URL that invokes a PL/SQL Agent to access the database. Here is the URL in question:

```
http://www.nhl.com:8080/ows-bin/nhl/owa/hockey_pass?person=Gretzky
```

1. The substring http://www.nhl.com:8080 in the above URL signals the Web browser to connect to the www.nhl.com host's port 8080 using the

HTTP protocol.

2. When the Oracle Web Listener that is running on www.nhl.com receives the request, the substring /ows-bin/nhl/owa signals the Web Listener to connect to the PL/SQL Agent instead of returning a file to the browser as it normally would have done with a static HTML document.

3. The URL is processed as follows:

   The Listener is configured so that the string /ows-bin/nhl/owa causes it to invoke the WRB Dispatcher. The WRB Dispatcher understands "owa" to mean it should pass execution to a WRBX interfaced to the PL/SQL Service. It passes the WRBX the following information in the form of CGI environment variables (For the sake of compatibility, WRB understands CGI environment variables):

   - The string /ows-bin/nhl/owa is passed as the environment variable SCRIPT_NAME. The PL/SQL Agent parses the SCRIPT_NAME to extract the DCD, which is nhl The DCD is passed as the directory name immediately preceding /owa, but is actually a file containing database connection information; owa is the name of the PL/SQL Agent itself. The use of DCDs is specific to the PL/SQL Agent.
   - The string /hockey_pass is passed as the environment variable PATH_INFO. This indicates that hockey_pass is the specific application to be executed. This application is a PL/SQL procedure stored in the database, that the PL/SQL Agent executes.
   - The string "person=Gretzky" is passed as the environment variable QUERY_STRING. This indicates that the value "Gretzky" is to be passed to the hockey_pass application for the parameter person. This will correspond in name and be compatible in datatype to a PL/SQL parameter used in the application.

The PL/SQL Agent

The WebServer employs the PL/SQL Agent to execute PL/SQL procedures stored in the database. If you install the PL/SQL Web Toolkit with your PL/SQL Agent, you can use a set of predefined PL/SQL packages to generate HTML formatted output, leaving you free to focus on the logic of your application. Even using the PL/SQL Web Toolkit, you need a conceptual understanding of HTML. For example, you must know at what points in your page anchors are necessary, although you needn't write the actual code for the anchors. The PL/SQL Agent takes care of interfacing to the needed environment variables and generating the HTML code.

Database Connection Descriptors (DCDs)

Whenever a URL invokes the PL/SQL Agent, it specifies a DCD. Creating and maintaining the DCDs is the responsibility of the WebServer administrator. As an application developer, all you need be concerned with is generating URLs that specify the correct DCD to achieve the result you want. The following is the information the DCD provides:

- *username*
- *password*
- *ORACLE_HOME*
- *ORACLE_SID*
- *SQL*Net V2 Service Name or Connect String*
- *owa_err_page*
- *owa_valid_ports*
- *owa_log_dir*
- *owa_nls_lang*

username

All SQL and PL/SQL statements are executed by the database under the auspices of some database user. This name identifies that user. The username determines the schema (logical section of a database) that the PL/SQL Agent connects to, the actions it can perform, the resources (disk space and so on) it can use, and the level of monitoring of its activities that the database performs.

password

This is provided if you want to password protect the DCD, ensuring that only certain people connect as the user identified by username. Keep in mind, however, that the users who connect to the database as username still will be able to do only what the application gives them the ability to specify through URLs.

ORACLE_HOME

This is the root of the Oracle7 Server code tree in the OS file system.

ORACLE_SID

If the PL/SQL Agent is to connect to a local database, this determines that database. If the WRB is used, the request will be handed off to a process connected already to this database, if possible

*SQL*Net V2 Service Name or Connect String*

If the PL/SQL Agent is to connect to a remote database, this determines that database.

owa_err_page

This is the pathname of the HTML document that the PL/SQL Agent is to return to the client's browser when an error occurs in the PL/SQL procedure that the PL/SQL Agent invoked. This is the actual path as understood by the OS, not the virtual path as configured in the Web Listener

owa_valid_ports

The valid Web Listener network ports to which the PL/SQL Agent will respond. The network port is, of course, specified in the URL. You must make sure that the URLs your applications generate provide port numbers that are valid for the WRB Service or CGI program they specify.

owa_log_dir

The directory where the PL/SQL Agent writes its error file. The error file differs from the page in that it is for the WebServer administrator's attention, rather than the clients.

owa_nls_lang

The NLS_LANG of the Oracle7 database to which the PL/SQL Agent connects. If not specified, the PL/SQL Agent administration program looks up the database NLS_LANG when the service is submitted. The NLS_LANG indicates the language to be used.

How the PL/SQL Agent Uses Environment Variables

Note: These are CGI standard environment variables. However, the PL/SQL Agent can use them whether it is invoked through *CGI* or through the Web Request Broker (WRB).

The PL/SQL Agent uses the environment variables shown below:

CGI Variables Used by the Oracle PL/SQL Agent

Variable	Contains
REQUEST_METHOD	GET or POST
PATH_INFO	the name of PL/SQL procedure to invoke

Variable	Contains
SCRIPT_NAME	contains the DCD the PL/SQL Agent is to use when logging on to Oracle7
QUERY_STRING	parameters to the PL/SQL procedure (for GET method only. POST method parameters are passed via standard input.)

Passing Parameters to PL/SQL

A PL/SQL procedure usually requires parameters in order to execute and generate the appropriate HTML document. The following section discusses several key concepts and tips that a PL/SQL developer should understand with respect to how parameters get passed to the specified PL/SQL routine.

These key concepts and tips are:

- *Getting Parameters from the Web Browser to the PL/SQL Agent*
- *Passing Parameters Using an HTML Form*
- *Providing Default Parameter Values*
- *Multivalued Parameters*
- *Overloading Procedures*

Getting Parameters from the Web Browser to the PL/SQL Agent

Depending on the REQUEST_METHOD used, parameters are passed from the Web Browser to the Web Listener to the PL/SQL Agent in one of two ways:

- Through the QUERY_STRING environment variable. If the GET method is used by the Web browser, the Web Listener passes the parameters to the PL/SQL Agent in this environment variable.
- Through standard input. If the POST method is used, the Web Listener passes the parameters to the PL/SQL Agent using standard input.

It is transparent to the PL/SQL procedure that is the actual consumer of these parameter(s) which method is used to pass the parameters from the Web Listener to the PL/SQL Agent and which protocol, the WRB or CGI, is used. This is an important feature of the Oracle PL/SQL Agent: the PL/SQL programmer need not be aware of whether GET or POST is used and need not be concerned with parsing either the QUERY_STRING environment variable or standard input. Thus, the PL/SQL programmer can concentrate on what he or she knows

best: developing the logic to extract data from the Oracle database, based on pre-parsed parameters passed by the Oracle PL/SQL Agent.

It is recommended that you use POST whenever possible. GET is the method used for links and non-form URLs. For HTML forms, one has a choice. Because the GET method uses operating system environment variables, there are limits on the length of the QUERY_STRING.

Passing Parameters Using an HTML Form

The following example is analogous to the one in the previous section, except that it uses an HTML form that employs the POST REQUEST_METHOD.

```
<FORM METHOD="POST" ACTION="http://www.nhl.com:8080/ows-bin/nhl/owa/
hockey_pass">
```

Please type the name of the person you wish to search for:

```
<INPUT TYPE="text" NAME="person"><P> To submit the query, press this
button: <INPUT TYPE="submit" VALUE="Submit Query">. <P> </FORM>
```

The above form does the same thing as the previous example, except that instead of populating the QUERY_STRING environment variable with "person=Gretzky," the Web Listener writes "person=Gretzky" to standard input.

Note that the name of the HTML input variable, in this case "person", has to be the same as the PL/SQL parameter it is to match.

The PL/SQL procedure that is the recipient of the above parameters follows:

```
create or replace procedure hockey_pass (person in varchar2) is
     n_assists integer;
     begin
     select num_assists into n_assists
       from hockey_stats
     where name=person;
htp.print(person||' has '||to_char(n_assists)||' assists this season')
end;
```

Providing Default Parameter Values

If you cannot guarantee that a value will be passed from a Web Browser for a particular PL/SQL procedure parameter, then you should give the parameter a default value. For example:

```
create or replace procedure showvals(a in varchar2 DEFAULT NULL,
                                     b in varchar2 DEFAULT NULL) is
begin
    htp.print('a = '||a||htp.br);
    htp.print('b = '||b||htp.br);
 end;
```

Suppose the PL/SQL Agent receives a request to call procedure showvals with no value for "a" and the value of 'Hello' for "b", and there was no DEFAULT

NULL clause in the procedure's definition. Then the request would generate an error with the following message:

```
OWS-05111: Agent : no procedure matches this call
   OWA_SERVICE: test_service
   PROCEDURE: showvals

  PARAMETERS:   ===========   B:    Hello
```

By "defaulting" the parameters, the above request would properly output:

```
a = <BR>
b = Hello<BR>
```

which to the end user would look like:

```
a =
b = Hello
```

Multivalued Parameters

Generally, you need not be concerned with the order in which the Oracle PL/SQL Agent receives parameters from a URL. The only case where it might be relevant is when passing multiple values for the same form field. In this case, all the values for that form field should be together, and if the order of the values in that field is significant, that order must be preserved in the URL.

There are a number of instances where you can have such multiple values for the same HTML variable. The HTML tag "SELECT" allows users to select from a set of possible values for an HTML form field. If the SIZE parameter of that SELECT tag is greater than one, the form allows multiple selection, and the user can choose to select more than one value. For example, if you ask a user to indicate her hobbies, she may well have more than one. In this case, you pass the multiple values to a single PL/SQL parameter, which must be a PL/SQL table. A PL/SQL table is a data structure similar to an array.

Example of a Multivalued Field

Another case where one has a set of values corresponding to a single form field is shown in this example:

QUERY_FORM prints an HTML page with all the columns for the specified table. Invoke the procedure from a Web Browser with a URL like: http://yourhost:port_num/service_name/owa/query_form?the_table=emp

```
create or replace procedure query_form(the_table in varchar2) is
     cursor cols is
    select column_name
          from user_tab_columns
    where table_name = upper(the_table);
 begin
    htp.htmlOpen;
    htp.headOpen;
    htp.htitle('Query the '||the_table||' table!');
```


ORCL000830

```
    htp.headClose;
    htp.bodyOpen;
    -- Use owa_util.get_owa_service path to automatically retrieve
    htp.formOpen(owa_util.get_owa_service_path||'do_query');
    -- Put in the table as a hidden field to pass on to do_query
    htp.formHidden('the_table', the_table);

    -- Put in a dummy value, as we cannot DEFAULT NULL a PL/SQL table.
    htp.formHidden('COLS', 'dummy');
    for crec in cols loop
       -- Create a checkbox for each column. The form field name
       -- will be COLS and the value will be the given column name
       -- Will need to use a PL/SQL table to retrieve a set of
       -- values like this. Can use the owa_util.ident_arr type
          -- since the columns are identifiers.
       htp.formCheckbox('COLS',crec.column_name);

       htp.print(crec.column_name);
       htp.nl;
    end loop;
    -- Pass a NULL field name for the Submit field; that way, a
    -- name/value pair is  not sent in.  Wouldn't want to do this
     -- if there were multiple submit buttons.
    htp.formSubmit(NULL, 'Execute Query');
       htp.formClose;
    htp.bodyClose;
    htp.htmlClose;
end;
```

Invoking this procedure brings up a page that looks like this:

Query the emp table!

- ☐ EMPNO
- ☐ ENAME
- ☐ JOB
- ☐ MGR
- ☐ HIREDATE
- ☐ SAL
- ☐ COMM
- ☐ DEPTNO

Execute Query

In this example, the user has already selected to query the EMPNO, ENAME, JOB, and SAL columns:

Here is a procedure to process this form submission:

```
-- DO_QUERY executes the query on the specified columns and
 -- tables.The OWA_UTIL.IDENT_ARR datatype is defined as:
 --   -- type ident_arr is table of varchar2(30) index by binary_integer
  create or replace procedure do_query(the_table in varchar2,
                                       cols in owa_util.ident_arr) is
    column_list varchar2(32000);
    col_counter integer;
    ignore      boolean;
begin
 -- For PL/SQL tables, have to just loop through until you hit
 -- no_data_found. Start the counter at 2 since we put in
 -- a dummy hidden field.
    col_counter := 2;
   loop
       -- build a comma-delimited list of columns
      column_list := column_list||cols(col_counter)||',';
       col_counter := col_counter + 1;
    end loop;
exception
    when no_data_found
    then
        -- strip out the last trailing comma
        column_list := substr(column_list,1,length(column_list)-1);
         -- print the table - assumes HTML table support
            ignore :=
owa_util.tablePrint(the_table,'BORDER',
OWA_UTIL.HTML_TABLE,  column_list);
 end;
```

Then, after selecting the "Execute Query" button, the user would see the following:

EMPNO	ENAME	JOB	SAL
7369	SMITH	CLERK	800
7499	ALLEN	SALESMAN	1600
7521	WARD	SALESMAN	1250
7566	JONES	MANAGER	2975
7654	MARTIN	SALESMAN	1250
7698	BLAKE	MANAGER	2850
7782	CLARK	MANAGER	2450
7788	SCOTT	ANALYST	3000
7839	KING	PRESIDENT	5000
7844	TURNER	SALESMAN	1500
7876	ADAMS	CLERK	1100
7900	JAMES	CLERK	950
7902	FORD	ANALYST	3000
7934	MILLER	CLERK	1300

If you cannot guarantee that at least one value will be submitted for the PL/SQL table, it is a good idea to use a hidden place-holder variable as the first value. The reason is that you cannot provide a default NULL value for a PL/SQL table, and a call to this procedure with just one argument (the_table) would cause the PL/SQL Agent to generate an error.

Note that the PL/SQL Agent can only pass parameters to PL/SQL tables that have a base type of VARCHAR2. This should not provide a significant limitation, as the PL/SQL type VARCHAR2 is the largest PL/SQL datatype, with a maximum length of 32K (32767 bytes). The values can then be explicitly converted to NUMBER, DATE, or LONG within the stored procedure (using TO_NUMBER or TO_DATE - no conversion needed for LONGs).

Overloading Procedures

As explained in the *Overview of the Oracle7 Server, SQL, and PL/SQL* under *Overloading Subprograms*, PL/SQL allows you to overload procedures and functions that are in PL/SQL packages. In overloading, multiple procedures or

functions have the same name, but are distinguished by the fact that they take different parameters. For example:

```
 create or replace package overload is
      procedure proc1(charval in varchar2);
      procedure proc1(numval  in number);
 end;
create or replace package body overload is
      procedure proc1(charval in varchar2) is
      begin
         htp.print('The character value is '||charval);
      end;
       procedure proc1(numval  in number);
         htp.print('The number value is '||numval);
end;
end;
```

Note: The PL/SQL Agent can use this functionality with the following restriction: if two procedures take the same number of parameters, those parameters must differ in name as well as datatype (normally a difference in datatype suffices to distinguish the parameters). For example:

```
create or replace package overload is
      procedure proc1(val in varchar2);
      procedure proc1(val in number);
 end;
```

If executed directly from SQL, this would be acceptable, but when the PL/SQL Agent attempts to determine which procedure to call, it is not able to distinguish between the two and will generate an error.

This limitation is imposed by the lack of HTML form datatypes

Oracle PL/SQL Agent Error Handling

There are two types of errors that the Oracle PL/SQL Agent handles:

- *Application Errors*. These are the responsibility of the application developer.
- *System Errors*. These are usually the responsibility of the WebServer administrator, although they also occur when application errors are not handled.

Application Errors

Application errors are specific to the PL/SQL application. All PL/SQL procedures you write should have their own exception handling (see *Handling Exceptions*) that produces the appropriate output in HTML form.

Because the PL/SQL Agent does not read the HTML output to determine its content, and properly handled exceptions themselves generate HTML error messages, handled exceptions are transparent. As far as the PL/SQL Agent is

concerned, if the PL/SQL code generates HTML output, the operation was successful. The user will see whatever handled exception message is generated by the PL/SQL procedure.

System Errors

System errors are detected by the Oracle PL/SQL Agent itself. These are errors that occur when the PL/SQL Agent is unable to launch the PL/SQL procedure or when a PL/SQL exception is not handled by the stored procedure, causing the exception to be propagated back to the PL/SQL Agent as a system error. This causes a standard HTML error document to be returned to the browser.

For example, if the Oracle PL/SQL Agent cannot make a connection to the Oracle7 Server, the PL/SQL procedure cannot run and a system error occurs. The PL/SQL Agent then returns a default error message to the browser or returns a customized HTML error page (if one was previously configured as part of the DCD using the OWA_ERR_PAGE parameter).



The Java™ Interpreter

The Java Interpreter is a part of the Oracle WebServer that interprets and executes Java and sends the HTML output to the client's browser as a Web page.

Oracle WebServer provides a set of Java classes to enable you to access the database and to generate HTML dynamically using the Java Interpreter. These classes are called the Java Web Toolkit. Using these Java classes, you can make PL/SQL calls from within your Java application, effectively combining the strengths of the two languages. This toolkit also includes a type wrapper for Java applets, to make it easy for you to store applets in the database, retrieve them through the Java Interpreter, and send them to the client for execution embedded in a Web page. You can also store and retrieve Java applets for execution on the client using the PL/SQL Agent by simply treating the applets as data.

Generally speaking, you should use Java to handle multimedia operations and to interface to objects on the net, and you should use PL/SQL for interfacing to the database. Whether this means using the PL/SQL Agent or the Java Interpreter will largely be determined by the following:

- Whether you want to pull all of your data from the database (as the PL/SQL Agent does) or whether you also want to access other data services such as the OS file system.
- The balance of database-intensive to net or multimedia intensive operations in your application.

To use the Java Web Toolkit, import the following into your Java code:

- *oracle.html.** contains the objects for dynamic HTML generation.
- *oracle.rdbms.** contains the objects for database access.
- *oracle.plsql.** contains the objects for PL/SQL access.

Obviously, if you only want to perform some of these functions, you need only import some of these objects. For more information, see *Database Access from Java*, *Example of Java Database Access*, *Dynamic HTML from Java*, and *Java Dynamic HTML Examples*.

Database Access from Java

Using Java to call *PL/SQL* circumvents the *PL/SQL Agent*, and this has a number of consequences, among them the following:

- PL/SQL Agent processes invoked through the WRB (but not CGI) connect to the database automatically at startup. Java processes connect to the database when necessary. This gives a significant performance advantage to the PL/SQL Agent.
- The PL/SQL Agent uses *Database Connection Descriptors (DCDs)* to control the privileges an application runs under and the database schema to which it connects in a generalized and application-independent way. In Java, database connections are coded into the application.

Nonetheless, Java can do many things that PL/SQL cannot, like access local files and manipulate multimedia objects. If you need this functionality executed on the server, the Java Interpreter is the way to go.

Creating Package Wrappers

The Java Interpreter interfaces to the Oracle7 Server by running *PL/SQL packages* or standalone PL/SQL procedures and functions. Each package an application is to run must have a *package wrapper*, which is a Java class containing methods to call that package's procedures and functions. Standalone procedures and functions are all wrapped in a single wrapper. Once you have identified or created the PL/SQL packages your applications needs, you can create the package wrappers for them by running the *pl2java* utility as follows:

```
pl2java [flags] username/password[@connect-string] packagename...
```

This utility creates a wrapper class for each package given as an argument to the command. When your application is run, it creates an instance of this class to interface to the package. If you have standalone procedures or functions in your applications, run the *pl2java* utility without any package names, but using the class flag as explained below. This will create a single class wrapper for all the standalone procedures and functions you use.

Here are the component definitions:

flags

Options that control how the wrappers will be created. These are explained below.

username

The name of the Oracle database user that owns the packages.

password

The password for the Oracle user identified by *username*.

connect-string

The string that identifies the local or remote database where the packages

are located. For local databases, this is the Oracle SID, as described in the *Oracle7 Server Administrator's Guide*. For remote databases, this is the SQL*Net Connect String, as described in *Understanding SQL*Net*.

package name...

A list of all the packages that your application references in the schema identified by username. To wrap standalone procedures and functions you must omit this component and must use the class flag to name the class wrapper that will be created. You should not include the containing schemas in the package names. It is good practice to keep all the packages, procedures, and functions you want to use in one schema.

flags

All of the flags that *pl2java* uses are optional, except, under certain conditions, class. Here are the descriptions of the flags:

-help

Provides help information.

-d <dir>

Sets the directory where the wrapper classes will be stored. The default is the current directory.

-package <packagename>

Sets the Java package to which the wrapper classes belong.

-class <class>

Sets the Java class to which the wrappers will belong. If the *pl2java* utility is run against packages, this flag is optional. Java classes based on packages inherit by default the names of the packages they encapsulate. This flag can override that default, but it only applies to the first package named in the command. If the wrappers are being created for standalone procedures and functions, then this flag is mandatory, and all procedures and functions named in the command are grouped into the single class named by this flag.

The names of the classes follow the capitalization given in the command. Since PL/SQL is not case-sensitive, this capitalization need not follow that actually given in the PL/SQL code itself.

There are certain PL/SQL datatypes that the *pl2java* utility cannot encapsulate. These are shown below, along with the recommended substitutes, if any:

Disallowed PL/SQL Datatype	Substitute PL/SQL Datatype
POSITIVE	BINARY INTEGER
PL/SQL table of BINARY INTEGER, NATURAL or POSITIVE	PL/SQL table of NUMBER
PL/SQL table of LONG	PL/SQL table of CHAR or VARCHAR2
PL/SQL table of BOOLEAN	PL/SQL table of NUMBER, *treat 0 as false, 1 as true*
ROWID	*none*
MSLABEL	*none*
PL/SQL table of ROWID	*none*
PL/SQL table of MSLABEL	*none*

Application Structure

Your Java program will be a Java class with a public static method called "main", which takes an array of strings, and various methods that are part of the wrapper classes. All the methods in the *Session* class throw *ServerException*. This exception is triggered when a database error occurs during the execution of the PL/SQL procedure or function. When invoking these methods, you should handle the exception appropriately.

The steps your application must go through are as follows:

1. It must create an object of type *Session* to handle the database connection. All of the operations performed during this session must be called from within this object. When the object exits, the database connection is marked for termination by the Java garbage collector. Since connecting to the database takes time, try to minimize the number of connections by grouping into one *Session* object all of the operations that involve a given schema.
2. For each PL/SQL package used in the application, it must create one instance of the packagewrapper subtype created for that PL/SQL package by the *pl2java* utility.

3. For each parameter of a PL/SQL package, it must create an instance of the Java variable that matches that parameter. All classes that encapsulate PL/SQL values have *toString*() methods. Therefore, you can concatenate the PL/SQL values directly in Java strings using Java's "+" concatenation operation.
4. Unlike Java and most languages, but like SQL, PL/SQL uses *Nulls and Three-Valued Logic* to deal with missing information. NULLs should be dealt with differently than known values. When you access a NULL from Java, it throws the *NullValueException* runtime exception. Therefore, you should account for this whenever a PL/SQL value may be NULL. You can do this in either of the following two ways:
   - Using the *isNull* method to check directly whether the value is NULL.
   - Enclosing the operation in a "try {...} catch (...) {...}" block that traps the *NullValueException*.

You must invoke the *Session setProperty* method to set parameters such as ORACLE_HOME that the database session requires. For more information on this method, refer to the standard Java hypertext reference included with this document.

You should try to logoff from the database explicitly when the session is no longer needed. Although Java garbage collection disconnects the session when the *Session* object is exited, this does not necessarily happen right away. Java does not perform garbage collection until resources are low or the program idles. Therefore, it is better to disconnect the session to free up the database connection resource immediately.

Example of Java Database Access

This section outlines a Java program that uses a PL/SQL package to look up all employees from a database and generates a report in HTML. In our example, you have an "EMP" database table and an "*Employee*" PL/SQL package that looks up employee information from that table. This table and package are created by user "scott" with password "tiger" in database "HR_DB".

Here is the SQL and PL/SQL source code that defines the objects used by this example.

```
CREATE TABLE EMP (
  emp_nameVARCHAR2(30) NOT NULL,
  emp_numberNUMBER(10),
  emp_deptVARCHAR2(30) NOT NULL
);

CREATE OR REPLACE package Employee as
  type string_table is table of varchar2(30) index by binary_integer;
```

```
  type number_table is table of number(10) index by binary_integer;

  function count_employees(
    dept_name   in  varchar2
  ) return number;

  procedure list_employees(
    dept_name   in varchar2,
    employee_name   out   string_table,
    employee_no  out  number_table
  );

end;

CREATE OR REPLACE package body Employee as

  function count_employees(
    dept_name  in  varchar2
  ) return number as
    employee_count  number;
  begin

    select count(*)
    into   employee_count
    from   EMP
    where  EMP_DEPT = dept_name;

    return employee_count;

  end;

  procedure list_employees(
    dept_name         in varchar2,
    employee_name     out string_table,
    employee_no       out  number_table
  ) as
    inumber;
    cursor employee_rec(dept_name varchar2) is
    select EMP_NAME, EMP_NUMBER
    from   EMP
    where  EMP_DEPT = dept_name;
  begin

    i := 1;
    for employee in employee_rec(dept_name) loop
      employee_name(i) := employee.EMP_NAME;
      employee_no(i)   := employee.EMP_NUMBER;
      i := i + 1;
    end loop;

  end;

end;
```

For reference, the commands themselves are considered SQL, but the code inside the packages is PL/SQL. Since PL/SQL is a superset of SQL, you can think of it all as PL/SQL.

This example uses the *oracle.html* package to generate dynamic HTML. That package is documented under *Dynamic HTML from Java*.

The main program

In this example, the program is called *EmployeeReport*. This program will be saved in a file called *EmployeeReport.java*.

```
import oracle.html.*;

public class EmployeeReport {

  public static void main (String args[]) {

    HtmlHead hd = new HtmlHead("Employee Listing");
    HtmlBody bd = new HtmlBody();
    HtmlPage hp = new HtmlPage(hd, bd);
    hp.printHeader();
    hp.print();

  }
}
```

Generating the Java Wrapper Class for the *Employee* Package

To access the *Employee* package, you must generate a Java wrapper class for it. This encapsulates the package as a Java class. The functions and procedures in the package will appear as methods in the wrapper class under the same names.

PL/SQL datatypes are encapsulated by Java objects. For example, a PL/SQL VARCHAR2 string is encapsulated by a Java *PStringBuffer* object, and a PL/SQL number is encapsulated by a Java *PDouble* object.

To generate the Java class wrapper for *Employee* package, execute the following command at a command prompt:

```
pl2java scott/tiger@HR_DB Employee
```

This generates the class wrapper *Employee.class* in the current directory. The class will have 1 constructor and 2 methods for the *Employee* package:

```
public class Employee {

  public Employee(Session session) { ... }

  public PDouble count_employees(PStringBuffer dept_name)
    throws ServerException { ... }

  public void list_employees(PStringBuffer dept_name,
     PStringBuffer employee_name[],
     PDouble       employee_number[])
    throws ServerException { ... }
}
```

Notice that the constructor takes a *Session* object as a parameter to encapsulate the database connection. This is because each time a PL/SQL package, procedure, or function is accessed in a session, it is instantiated. In effect, a copy is created that is the private property of that session. This gives the package a stable state for that session, while leaving it free to have another state when called by another session.

Overriding Default Value Sizes

When a PL/SQL function returns a value whose size is variable - for example VARCHAR2, LONG, RAW, or LONG RAW - the size of the value is set by

default to 255 bytes. In the wrapper class, you may change the default size by setting the following variable for the PL/SQL function in question:

```
<function name>_<overload number>_return_length
```

The overload number is the number of other functions that exist with the same name as this one. To find out about overloading of functions in PL/SQL, see *Overloading Subprograms*. For more information on overload numbers specifically, see the *PL/SQL User's Guide and Reference*. You can find out what the overload number of a function is by using the Oracle7 Server standard package *dbms_describe*, as covered in the *Oracle7 Server Administrator's Guide*. For non-overloaded functions, the overload number is 0.

For example, assume the following PL/SQL package:

```
CREATE OR REPLACE package Employee as
  function employee_name (
    employee_numberinnumber
  ) return varchar2;

END;
```

By default, the length of the return value is 255 bytes. You can change this using the code shown below:

```
public class EmployeeReport {
  public static void main(String args[]) {
    Session session = new Session("scott", "tiger", "HR_DB");
    Employee employee = new Employee(session);
    PDouble pEmployeeNumber = new PDouble((double)12345);
    // Set VARCHAR2 return length for function employee_name to 50
employee.employee_name_0_return_length = 50;
    System.out.println("Employee ID: " + pEmployeeNumber +
       " name: " + employee.employee_name(pEmployeeNumber));    }
}
```

Similarly, when a PL/SQL function returns a PL/SQL table, its default length is 40. You can change that length by changing the following variable in the wrapper class:

```
<function name>_<overload number>_return_arraylength
```

Making a connection to the database

The first thing that you have to do in the *EmployeeReport* program is to connect to a database. You do this by using the following code to create a *Session* object:

```
import oracle.html.*;
import oracle.rdbms.*;// ADD: import Oracle classes which deal //with
database

public class EmployeeReport {

  public static void main (String args[]) {

    HtmlHead hd = new HtmlHead("Employee Listing");
    HtmlBody bd = new HtmlBody();
```

```
    HtmlPage hp = new HtmlPage(hd, bd);
    hp.printHeader();


    // ADD: defines Oracle session properties like ORACLE_HOME
    Session.setProperty("ORACLE_HOME", "/user/oracle");
    Session.setProperty("TNS_ADMIN", "/user/oracle/network/admin");

    // ADD: creates a database session and logon
    Session session;
    try {
      session = new Session("scott", "tiger", "HR_DB");
    } catch (ServerException e) {
      bd.addItem(new SimpleItem("Logon fails: " + e.getSqlerrm()));
      hp.print();
      return;
    }
    hp.print()
  }
}
```

To handle any errors raised, put the operation in a "try {...} catch (...) {... }" block and trap any *ServerExceptions*.

Invoking the *Employee* package

To invoke the *Employee* package, you need to create a new instance of the corresponding wrapper class. Then you can call the procedures and functions in the package by invoking the methods in the wrapper class. Add the following code to the program:

```
import oracle.html.*;
import oracle.rdbms.*;
import oracle.plsql.*;// ADD: import Oracle classes which deal //with PL/
SQL data types

public class EmployeeReport {

    public static void main (String args[]) {

    HtmlHead hd = new HtmlHead("Employee Listing");
    HtmlBody bd = new HtmlBody();
    HtmlPage hp = new HtmlPage(hd, bd);
    hp.printHeader();

    Session.setProperty("ORACLE_HOME", "/user/oracle");
    Session.setProperty("TNS_ADMIN", "/user/oracle/network/admin");

    Session session;
    try {
      session = new Session("scott", "tiger", "HR_DB");
    } catch (ServerException e) {
      bd.addItem(new SimpleItem("Logon fails: " + e.getSqlerrm()));
      hp.print();
      return;
    }

    // ADD: create a new instance of Employee package
    Employee employee = new Employee(session);

    // ADD: find the department name from the input parameter
    String deptName = null;
    if ((args.length < 1) || !args[0].startsWith("DEPT=")) {
      bd.addItem(new SimpleItem("No department name given"));
      hp.print();
```

```
      return;
    } else {
      deptName = args[0].substring(5);
    }

    // ADD: create objects to encapsulate PL/SQL values that are
    // used as parameters
    PStringBuffer pDeptName = new PStringBuffer(30, deptName);
    PStringBuffer pEmployeeName[];
    PDouble       pEmployeeNumber[];
    PDouble  pEmployeeCount;

    // ADD: print report header
    bd.addItem("Department " + pDeptName + ":")
      .addItem(SimpleItem.Paragraph);

    // ADD: call Employee package to count the number of employees in
    // the department
    try {
      pEmployeeCount = employee.count_employees(pDeptName);
    } catch (ServerException e) {
      bd.addItem("Fail to retrieve employee information for department " +
deptName + ": " + e.getSqlerrm());
      hp.print();
      return;
    }
    int employeeCount = (int)pEmployeeCount.doubleValue();
    if (employeeCount == 0) {
      bd.addItem("No employee found under department " + deptName);
      hp.print();
      return;
    }

    // ADD: allocate the arrays for employee names and numbers
    pEmployeeName     = new PStringBuffer[employeeCount];
    pEmployeeNumber   = new PDouble[employeeCount];

    // ADD: allocate the buffers to retrieve employee information
    for(int i = 0; i < employeeCount; i++) {
      // max length of employee name is 30 (characters)
      pEmployeeName[i]   = new PStringBuffer(30);
      pEmployeeNumber[i] = new PDouble();
    }

    // ADD: call Employee package to look up employees in the dept
    try {
      employee.list_employees(pDeptName, pEmployeeName, pEmployeeNumber);
    } catch (ServerException e) {
      bd.addItem("Fail to retrieve employee information for department " +
deptName + ": " + e.getSqlerrm());
      hp.print();
      return;
    }

    // ADD: generate report
    DynamicTable tab = new DynamicTable(2);
    TableRow row = new TableRow();
    row.addCell(new TableHeaderCell("Employee Name"))
       .addCell(new TableHeaderCell("Employee Number"));
    tab.addRow(row);

    for (int i = 0; i < employeeCount; i++) {
      row = new TableRow();
      if (pEmployeeNumber[i].isNull())
row.addCell(new TableDataCell(pEmployeeName[i].toString()))
   .addCell(new TableDataCell("new employee"));
      else
row.addCell(new TableDataCell(pEmployeeName[i].toString()))
   .addCell(new TableDataCell(pEmployeeNumber[i].toString()));
      tab.addRow(row);
```

```
    }
    hp.addItem(tab);
    hp.print();

    // ADD: logoff from database
    try {
      session.logoff();
    } catch (ServerException e);
  }
}
```

There are a few aspects of the preceding code worth pointing out:

1. You must create an instance of Employee package before you can invoke the procedures and functions in it. When initiating the package, you need to specify the database session where the instance of the package is to be created.
2. Before you invoke a PL/SQL procedure or function, you have to create Java instance variables for the PL/SQL values that are to be used as parameters (like "pDeptName" in the above example). PL/SQL tables map to Java arrays. Remember to allocate the array as well as the individual elements in the array.
3. When you access a PL/SQL value, remember that the value may be NULL unless database constraints prevent this. If the value is NULL, and you try to retrieve it using the value methods (such as doubleValue() of PDouble), it throws a *NullValueException* runtime condition. It is better to ensure the value is not NULL before retrieving the value or to try the *NullValueException.*
4. All classes that encapsulate PL/SQL values have *toString*() methods and therefore can be concatenated. For an example of this, see the report-header-generation section above.
5. You should try to logoff from the database explicitly when the session is no longer needed. When a session is no longer needed, the session is disconnected when Java performs garbage collection. Java does not, however, guarantee that any garbage objects will be collected immediately when they become garbage. In fact, Java's garbage collector waits until the program idles, which for a busy Web site could be infrequent, or until resources are low, before it collects garbage objects. Therefore, it is better to issue a disconnect statement to free up database resources explicitly.

Dynamic HTML from Java

To generate dynamic HTML from within Java, you create various objects that use the interface *IHtmlItem.* All classes that generate dynamic HTML implement this interface, which has two simple methods: *toHTML* and *print*. Both of these methods produce the content of the object as HTML, but *toHTML* returns it as a string, whereas *print* sends it to system output. Therefore, you use *toHTML* when

you want the resulting HTML to be further processed by another method and use *print* when you want to output it. In effect, you build your Web page in a buffer with *toHTML* and flush the buffer with *print*.

The *oracle.html* package provides a standard set of classes based on HTML2, HTML3, and popular browser-specific extensions. You are not limited to these, however. You can easily create your own customizable HTML classes by deriving them from the *CompoundItem* or *Container* classes. The *oracle.html* package also has the intelligence to generate output that is optimized for the browser at hand. For example, a browser that does not support tables will get table data in the form of preformatted strings.

In some cases, interfaces have been used to specify the attributes of HTML tags. This was done to simplify cases where tag assignments can be complex or where similar arguments are used by several types of tags.

If you have a body of HTML you want to use repeatedly, you can encapsulate it in an object of class *Compounditem* and thereafter treat it as a single *HTMLitem*.

The HTML tags that you can dynamically generate using the supplied objects are listed below:

HTML FEATURE	JAVA OBJECTS THAT GENERATE
headings	HTMLHead
page breaks	HTMLPage
main body of page	HTMLBody
Comments	Comment
links	Link
anchors	Anchor
client- side Java applets	Applet
checkboxes	CheckBox
forms	Form
lists	Java classes exist for various types
frames	Frameset; Frame
hidden fields	Hidden
GIF graphics	Image
select options	Select; Option
passwords	PasswordField
radio buttons	Radio
tables	Table; DynamicTable; TableRow; TableCell; TableDataCell; TableHeaderCell; TableRowCell
text areas	TextArea; TextField

As you can see, the Java objects that generate the main structural HTML tags begin with HTML; others are chiefly named for the tags they generate.

The general procedure is to use the first three objects to define the basic structure of your generated Web page to then to use the *AddItem* method to add *HTMLItem*s to the body.

Java Dynamic HTML Examples

Here are examples of how to dynamically generate some HTML text using the Java Interpreter.

The following is a basic Java program that produces an HTML page whose title and content are both the famous "Hello World!":

```
import oracle.html.*;

public class HelloWorld {

public static void main (String args[]) {

  // Create an HtmlHead Object titled "Hello World!"
  HtmlHead hd = new HtmlHead("Hello World!");

  // Create an HtmlBody Object
  HtmlBody bd = new HtmlBody();

  // Create an HtmlPage Object
  HtmlPage hp = new HtmlPage(hd, bd);

  // Adds a simple string "Hello World" in this page
  bd.addItem("Hello World!");

  // Print out the content of this Page
  hp.print();
}

}
```

The following Java code creates an HTML anchor:

```
// Creates an anchor
Anchor anchor = new Anchor("expire_date", new SimpleItem("Expire Date: 02/
96"));
```

The following Java code creates an HTML form:

```
// Create a form object
Form form = new Form("GET", "http://www.myhom.com/wrb/doit");

// Create a TextField object and add it to the form
form.addItem(new TextField("textfield"));

// Add the form object to the HtmlBody object
bd.addItem(form);
```

The following Java code creates an HTML table. To make this example realistic, we have added some user-defined functions:

```
// Some user-defined functions
Product product = getFirstProduct();

// create a dynamic table with 2 columns
DynamicTable tab = new DynamicTable(2);

// create the rows and add them to the table
TableRow rows[] = new TableRow[NUM_ROWS];

for (int i=0; i< NUM_ROWS; i++) {
  // allocate TableRow
  rows[i] = new TableRow();
  // populate row with data
  rows[i].addCell(new TableHeaderCell(product.getProductID()))
         .addCell(new TableDataCell(product.getProductDescription()))
```

```
  // add them to Table
  tab.addRow(rows[i]);
}
```

The following Java code creates an HTML menu:

```
// Create a MenuList Object
MenuList menulist = new MenuList();
// Add new items to the list
menulist.addItem(new SimpleItem("Menu Item 1"))
        .addItem(new SimpleItem("Menu Item 2"));
// Add the list object to the body
bd.addItem(menulist);
```

The following Java code creates an HTML definition list, encapsulated in a *Container* object:

```
// Creates a new Container Object
Container dterms = new Container();
dterms.addItem(new SimpleItem("DefTerm1.1"));
dterms.addItem(new SimpleItem("DefTerm1.2"));
DefinitionList dl = new DefinitionList();
// Creates a new Definition List Object, note the first argument
dl.addDef(dterms, new SimpleItem("Definition1"));
```

The following Java code creates an HTML ordered list:

```
// Create a OrderedList Object
OrderedList orderedlist = new OrderedList();
// Add new items to the list
orderedlist.addItem(new SimpleItem("Ordered Item 1"))
           .addItem(new SimpleItem("Ordered Item 2"));
// Add the list object to the body (assuming it already exists)
bd.addItem(orderedlist);
```

The following Java code encapsulates a group of HTML tags as a single component that can be used repeatedly in a page:

```
// Create a new CompoundItem object
CompoundItem compoundItem = new CompoundItem();
// Set the default text attribute of all items in Compound Item
// Note that this operation cannot be done with a <b>Container</b>
compoundItem.setItal();
// Add a SimpleItem and a TextArea to the CompoundItem
compoundItem.addItem(new Simpleitem("How are you?").setBold())
            .addItem(new TextArea("response", 30, 10));
// Add the CompoundItem to the body object (assuming it's been created)
bd.addItem(compoundItem);
```

The following Java code specifies an applet of Java bytecode to be included in the Web page output for execution on the client's browser.

```
// Create a new Applet object with the following attributes:
// Applet file name:  "NervousText.class"
// Width of Applet Window:  400
// Height of Applet Window:  75
// Parameter:   Name="text", Value="This is an applet test"
Applet applet = new Applet("NervousText.class", 400, 75);
```

```
applet.addParam("text", "This is an applet test.");
```

The LiveHTML Interpreter

Files that the WebServer is to parse for LiveHTML are internally of the following MIME type:

```
text/x-server-parsed-html
```

The WebServer Administrator normally will create file extensions that are synonyms for this type, and these are what you use in your application code. The default synonym is SHTML. The administrator can also specify HTML as a synonym, in which case all HTML files are parsed for LiveHTML. Unless all your HTML files actually use LiveHTML, this is a bad idea, as it degrades performance. The WebServer Administrator can also specify that some Web Listeners are to allow full LiveHTML, some crippled LiveHTML, and some no LiveHTML at all. An application developer needs to know which Listeners are which, so as to specify the correct port numbers in the URLs. You can find this out by examining the WebServer Manager

LiveHTML code is formatted as SGML comments, so that it is ignored should the file ever find its way to the browser unparsed. The format for LiveHTML codes, therefore, is the following:

```
<!--#command tag1="value1" tag2="value2" -->
```

The tags are arguments to the commands, most of which actually only accept one of the possible tags. The possible commands and their associated tags are as follows:

- *config.* This command sets parameters for how the file or script is to be parsed and therefore is normally the first LiveHTML command in a file. The possible tags are:
  - *errmsg.* This specifies the error message that is sent to the client if an error occurs while parsing this document. Here is an example:

```
<!--#config errmsg="A parse error occurred in the Hockey_Pass file"-->
```

  - *timefmt.* This specifies a date format. LiveHTML files frequently include timestamps. The conventions follow the *strftime* library call supported in most versions of Unix, even if the WebServer is not running on a Unix platform.

- *sizefmt.* This specifies the format used when displaying a file size. The possibilities are *bytes*, which gives the absolute size in bytes, and *abbrev*, which gives the size in kilobytes or megabytes as appropriate.
- *cmdecho.* This specifies whether non-*CGI* scripts subsequently executed have their output incorporated into this HTML page. The possible values are *ON* and *OFF*. ON specifies that the output is included. The default is OFF.
- *cmdprefix.* Specifies a string that will be prepended to each line of the script output.
- *cmdpostfix.* Specifies a string that will be appended to each line of the script output.

- *include.* This command specifies that a file is to be included in the generated HTML page at this point. The file can any of the following:
  - another LiveHTML file like the current one.
  - a regular HTML file.
  - an ASCII file.

  which of these it is determined, as usual, by the extension. The possible tags are:
  - *virtual.* This gives a *virtual path* to the file. The directory mappings for virtual paths are set by the WebServer administrator using WebServer Manager.
  - *file.* This gives a pathname relative to the current directory. References to parent directories or uses of absolute pathnames are forbidden.
- *echo.* This gives the value of an environment variable. This variable is either one of the standard CGI environment variables or one of the LiveHTML extensions, which currently are all standard Server Side Include variables. There is only one tag, *var*, and it must be present. It provides the name of the variable. The LiveHTML environment variables are as follows:
  - *DOCUMENT_NAME.* The current filename.
  - *DOCUMENT_URL.* The virtual path to this file.
  - *QUERY_STRING_UNESCAPED.* If the client sent a query string, this is an unescaped version of it, with all shell-special characters escaped with \.
  - *DATE_LOCAL* The current date and local time zone in the format specified by the most recent *config timefmt* command.
  - *DATE_GMT.* Same as the above, but in Greenwich mean time.

  - *LAST_MODIFIED.* The last modification date of the file, given in the format specified in the last *config timefmt* command.

- *fsize.* This produces the size of the file in the format specified in the most recent *config filesize* command. Tags are the same as for *include.*
- *flastmod.* This produces the last modification date of the file in the format specified in the most recent *config timefmt* command. Tags are the same as for *include.*
- *exec.* This is the command to execute a script. The tags specify whether or not the script is CGI.
  - *cmd.* This specifies a non-CGI script. Execution is passed to the Operating System, and the given string is parsed as though it were entered at a command-line interface. The full path of the script must be given. The non-CGI environment variables specified under *echo* above can be referenced. Whether the output of the script is included in the HTML page that the parser outputs, is determined by the most recently executed *config cmdencho.*
  - *cgi.* This specifies a CGI script. The value given will be the *virtual path* of the script. *URL* locations are automatically converted into HTML anchors.

CHAPTER

5 Sample Applications

This section provides an overview of the simple applications included with the Oracle WebServer. These applications fall into the following categories:

- *PL/SQL Agent Sample Applications*
- *LiveHTML Sample Tags*
- *Java Sample Applications*

PL/SQL Agent Sample Applications

There are three sample applications using the PL/SQL Agent:

- Hello World. This creates a simple Web page incorporating an environment variable.
- Variable Field Forms. This generates varying text depending on which boxes the user clicks in an HTML form.
- Cookie. This uses a seven-day cookie to tell the user how often he has visited this sight.

LiveHTML Sample Tags

This provides examples of how to do the following with LiveHTML tags:

- Include files from the OS.
- Reference environment variables.

- Set Data and Time formats.
- Reference true or virtual document paths.
- Reference sizes and modification dates of OS files.
- Reference information about the WebServer.
- Reference information about the user's browser.

Java Sample Applications

There are two sample applications using the Java Interpreter:

- Hello World. This simply outputs a string.
- Advanced Java Applications. This shows how to generate a number of HTML tags dynamically.

CHAPTER

6

The PL/SQL Web Toolkit Reference

This section describes the hypertext procedures, hypertext functions, and utilities that make up the Oracle WebServer PL/SQL Web Toolkit.

One of the main goals of the Oracle PL/SQL Agent is to eliminate the PL/SQL programmer's need to be intimately familiar with World Wide Web technology. To this end, the Oracle WebServer includes a PL/SQL Web Toolkit made up of several PL/SQL packages that minimize the programmer's need to know HTML syntax. Although the programmer is still required to have a working knowledge of HTML, by using the Toolkit he or she will not need to hard code the exact syntax of HTML tags into PL/SQL procedures. For instance, a programmer still needs to realize that an anchor tag is needed, but he or she doesn't need to know the exact sequence of characters needed to generate an anchor.

The PL/SQL Web Toolkit includes the following PL/SQL packages:

Hypertext Procedures (HTP)

A hypertext procedure generates a line in an HTML document that contains the HTML tag that corresponds to its name. For instance, the htp.anchor procedure generates an anchor tag. The HTP package will be the most commonly used package of the three.

Hypertext Functions (HTF)

A hypertext function returns the HTML tag that corresponds to its name. However, it is not sufficient to call an HTF function on its own because the HTML tag is not passed to the PL/SQL Agent. The output of an HTF function must be passed to htp.print in order to actually be part of an HTML document. Thus, the following line:

```
htp.print(htf.italic('Title'));
```

is functionally equivalent to:

```
htp.italic('Title');
```

Every hypertext function (HTF) has a corresponding hypertext procedure (HTP). Thus, HTF functions are generally used only when the programmer needs to nest calls, such as:

```
htp.header(1,htf.italic('Title'));
```

In this example, 'htf.italic' will generate the following character string:

```
<I>Title</I>
```

This string is then passed to the 'htp.header' procedure and the following line will appear in the HTML document being formatted:

```
<H1><I>Title</I></H1>
```

OWA Utilities (OWA_UTIL)

This is a collection of useful utility procedures and functions. The purposes of these range from printing a signature tag on HTML pages to easy formatting of Oracle tables into HTML.

OWA

This is a set of procedures called only by the Oracle PL/SQL Agent itself. None of the subprograms in this package should be called directly from user-developed PL/SQL.

Pattern Matching Utilities (OWA_PATTERN)

This is a set of procedures and functions you can use to perform string matching and substitution with rich regular expression functionality.

Text Manipulation Utilities (OWA_TEXT)

This is a set of procedures, functions, and datatypes used by OWA_PATTERN for manipulating large data strings. They are externalized so you can use them directly if you wish.

Image Map Utilities (OWA_IMAGE)

This is a set of datatypes and functions for manipulating HTML image maps.

Cookie Utilities (OWA_COOKIE)

This is a set of datatypes procedures, and functions for manipulating HTML cookies.

Installing the Oracle WebServer PL/SQL Web Toolkit

To install the PL/SQL Web Toolkit, use the Oracle PL/SQL Agent DCD Administration forms. Both the DCD Creation and DCD Modification pages provide a checkbox for installing the PL/SQL Web Toolkit. Selecting this button and submitting the form will do the following:

- grant the CONNECT and RESOURCE roles to the OWA database user For more information on these roles or on granting roles, see "GRANT (roles)" in Chapter 4 of the *Oracle7 Server SQL Reference*.
- execute the OWAINS.SQL sql script, which can be found in the ows Administration directory. If run manually, the script should be run from SQL*DBA or from Server Manager. If you want to run it from SQL*Plus, see the header of the script for instructions.

  The OWAINS.SQL script installs all of the PL/SQL Web Toolkit packages.

Optimizing Multiple-DCD Installations

If your site has multiple PL/SQL Agent DCDs, you can minimize the amount of storage space used and enhance PL/SQL performance by doing the following:

1. Install the Pl/SQL Web Toolkit in one database user's schema. This user becomes the toolkit owner.
2. Drop the Pl/SQL Web Toolkit PL/SQL from the schemas of the OWA database users for other PL/SQL Agent DCDs, if you have already installed them.

```
connect <user> / <password>
drop package HTF;
drop package HTP;
drop package OWA_UTIL;
drop package OWA;
drop package OWA_PATTERN;
drop package OWA_TEXT;
drop package OWA_IMAGE;
drop package OWA_COOKIE;
```

```
drop package OWA_INIT;
```

3. Grant the system privilege EXECUTE on the PL/SQL packages to OWA database users for other PL/SQL Agent DCDs. For more information on this command, see "GRANT (system privileges)" in Chapter 4 of the *Oracle7 Server SQL Reference*.

```
connect <toolkit owner> / <password>
grant execute on HTF to <user>;
grant execute on HTP to <user>;
grant execute on OWA_UTIL to <user>;
grant execute on OWA to <user>;
grant execute on OWA_PATTERN to <user>;
grant execute on OWA_TEXT to <user>;
grant execute on OWA_IMAGE to <user>;
grant execute on OWA_COOKIE to <user>;
grant execute on OWA_INIT to <user>;
```

4. Create synonyms for the PI/SQL Web Toolkit PL/SQL packages in the schemas of all OWA database users who are specified in PL/SQL Agent DCDs. Synonyms are alternate names that make it possible for the packages to be referred to without being qualified by schema names. For more information, see "Ownership and Naming Conventions" in Chapter of this manual and see "CREATE SYNONYM" in Chapter 4 of the *Oracle7 Server SQL Reference*.

```
connect <user>/<password>
create synonym HTF for <Toolkit owner>.HTF;
create synonym HTP for <Toolkit owner>.HTP;
create synonym OWA_UTIL for <Toolkit owner>.OWA_UTIL;
create synonym OWA for <Toolkit owner>.OWA;
create synonym OWA_PATTERN for <Toolkit owner>.OWA_PATTERN;
create synonym OWA_TEXT for <Toolkit owner>.OWA_TEXT;
create synonym OWA_IMAGE for <Toolkit owner>.OWA_IMAGE;
create synonym OWA_COOKIE for <Toolkit owner>.OWA_COOKIE;
create synonym OWA_INIT for <Toolkit owner>.OWA_INIT;
```

Security Note

PL/SQL procedures run with the privileges of the creator of the PL/SQL code. For the Web Toolkit, this is only an issue for the owa_util package. Two of the subprograms, showsource and tableprint, access user data. Granting execute privileges on this package to users allows those users to view the tables, views, and stored PL/SQL code of the owner of owa_util.

If this is a security issue for your installation, install the owa_util package separately for each OWA database user.

The scripts to do this are:

- PUBUTIL.SQL
- PRIVUTIL.SQL

Both scripts reside in the OWS Administration directory, and should be run in the order listed.

Procedure and Function Reference

This section describes each procedure and function in the htp, htf, and owa_util packages. Please note that for every htp procedure that generates HTML tags, a corresponding htf function exists with identical parameters. Note that defaulted parameters do not need to be passed.

The description of each procedure or function is broken down into the following parts. Note that items for which no "generates" entry is shown are available only as procedures.

Parameters Passed into Procedures and Functions

All parameters passed into a hypertext procedure or function are of data type varchar2 (varying-length character string), integer, or date. The data type is indicated by the first letter of the parameter's name, "c" for character (varchar2), "n" for number (integer), and "d" for date. For example:

```
cname in varchar2
```

The "c" in cname indicates a character data type (varchar2).

```
nsize in integer
```

The "n" in nsize indicates a number data type (integer).

```
dbuf in date
```

The "d" in dbuf indicates a date data type (date).

As in BNF notation, a vertical bar (|) in the syntax diagram means "or".

Note: Many HTML 3.0 tags have a large number of optional attributes that, if passed as individual parameters to the hypertext procedures or functions would make the calls quite cumbersome. In addition, some browsers support non-standard attributes. Therefore, each hypertext procedure or function that generates an HTML tag has as its last parameter *cattributes*, an optional parameter. This parameter enables you to pass the exact text of the desired HTML attributes to the PL/SQL procedure.

For example, the syntax for htp.em is:

```
htp.em (ctext, cattributes);
```

A call that uses HTML 3.0 attributes might look like the following:

```
htp.em('This is an example','ID="SGML_ID" LANG="en"');
```

This line would generate the following:

```
<EM ID="SGML_ID" LANG="en">This is an example</EM>
```

Print Procedures

The following print procedures are used in conjunction with htf functions to generate a line in the HTML document being constructed. They can also be passed hard-coded text that will appear in the HTML document as-is. The generated line is passed to the PL/SQL Agent, which sends it to standard output. As documented in the CGI 1.1 specification, the Oracle Web Listener takes the contents of standard output and returns it to the Web browser that requested the dynamic HTML document.

htp.print

Syntax

htp.print (cbuf | dbuf | nbuf);

Purpose

generates a line in an HTML document.

Parameters

cbuf in varchar2 *or*
dbuf in date *or*
nbuf in number

Generates

Generates a line in an HTML document based on the value passed to it.

htp.prn

Alias for htp.print

Syntax

htp.prn (cbuf | dbuf | nbuf);

Purpose

Just like htp.print, but doesn't put a new line at the end of the value submitted.

htp.prints

Syntax

htp.prints (ctext);

Purpose

Generates a line in an HTML document and replaces all occurrences of the following special characters with the shown escape characters. If not replaced, the special characters would be interpreted as HTML control characters, and would produce garbled output.

'<' with '<'
'>' with '>'
'"' with '"'
'&' with '&'

Parameters

ctext in varchar2

Generates

Generates a line in an HTML document based on the value passed to it. This procedure is the same as htp.print or htp.p but first replaces the special characters listed above with escape characters.

htp.ps

Alias for htp.prints

Structure Tags

The following tags are used to identify the major parts of an HTML document.

Note: Although this section shows hypertext procedures (HTP), all of them are also available as hypertext functions (HTF).

htp.htmlOpen

Syntax

htp.htmlOpen;

Purpose

Prints a tag that indicates the beginning of an HTML document

Parameters

none

Generates

```
<HTML>
```

htp.htmlClose

Syntax

htp.htmlClose;

Purpose

Prints a tag that indicates the end of an HTML document

Parameters

none

Generates

```
</HTML>
```

htp.headOpen

Syntax

htp.headOpen;

Purpose

Prints a tag that indicates the beginning of the HTML document head

Parameters

none

Generates

```
<HEAD>
```

htp.headClose

Syntax

htp.headClose;

Purpose

Prints a tag that indicates the end of the HTML document head

Parameters

none

Generates

```
</HEAD>
```

htp.bodyOpen

Syntax

htp.bodyOpen (cbackground, cattributes);

Purpose

Prints the tag that identifies the beginning of the body of an HTML document, and allows you to specify an image as the background of the document

Parameters

cbackground	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<BODY background="cbackground" cattributes>
```

Note: If cbackground and cattributes are NULL, this tag generates `<BODY>`.

Example

```
htp.bodyOpen ('/img/background.gif');
```

This line produces:

```
<BODY background="/img/background.gif">
```

htp.bodyClose

Syntax

htp.bodyClose;

Purpose

Defines the end of the HTML document body

Parameters

none

ORCL000864

Generates

```
</BODY>
```

Head Related Tags

The following procedure tags should be placed between the htp.headOpen and htp.headClose procedure tags.

Note: Although this section shows hypertext procedures (HTP), all of them are also available as hypertext functions (HTF).

htp.title

Syntax

htp.title (ctitle);

Purpose

Prints an HTML tag with the text you pass in as the value of TITLE. Most Web Browsers display the text value enclosed between <TITLE> and </TITLE> at the top of the document viewing window.

Parameters

ctitle in varchar2

Generates

```
<TITLE>ctitle</TITLE>
```

htp.base

Syntax

htp.base (ctarget, cattributes);

Purpose

Prints an HTML tag that records the URL of the document

Parameters

ctarget	in varchar2 DEFAULT NULL
cattributest	in varchar2 DEFAULT NULL

Generates

Inserts absolute pathname of current document.

htp.isindex

Syntax

htp.isindex (cprompt, curl);

Purpose

Creates a single entry field with a prompting text, such as "*enter value,*" then sends that value to the URL of the page or program.

Parameters

cprompt	in varchar2	DEFAULT NULL
curl	in varchar2	DEFAULT NULL

Generates

```
<ISINDEX PROMPT="cprompt" HREF="curl">
```

htp.linkRel

Syntax

htp.linkRel (crel, curl, ctitle);

Purpose

Prints the HTML tag that gives the relationship described by the hypertext link from the anchor to the target. This is only used when the HREF attribute is present. This tag indicates a relationship between documents, but does not create a link. To do that, use htp.anchor.

Parameters

crel	in varchar2	
curl	in varchar2	
ctitle	in varchar2	DEFAULT NULL

Generates

```
<LINK REL="crel" HREF="curl" TITLE="ctitle">
```

htp.linkRev

Syntax

htp.linkRev (crev, curl, ctitle);

Purpose

Gives the relationship described by the hypertext link from the target to the anchor. This is the opposite of htp.linkRel. This tag indicates a relationship between documents, but does not create a link. To do that, use htp.anchor.

Parameters

crev	in varchar2	
curl	in varchar2	
ctitle	in varchar2	DEFAULT NULL

Generates

```
<LINK REV="crev" HREF="curl" TITLE="ctitle">
```

htp.meta

Syntax

htp.meta (chttp_equiv, cname, ccontent);

Purpose

Prints an HTML tag that identifies and embeds document meta-information that supplies the Web browser with information about the objects returned in HTTP.

Parameters

chttp_equiv	in varchar2
cname	in varchar2
ccontent	in varchar2

Generates

```
<META HTTP-EQUIV="chttp_equiv" NAME ="cname" CONTENT="ccontent">
```

Example

```
htp.meta ('Refresh', NULL, 120);
```

This line produces:

```
<META HTTP-EQUIV="Refresh" CONTENT=120>
```

which on some Web browsers will cause the current URL to be reloaded automatically every 120 seconds.

Body Tags

Body tags are used in the main text of your HTML page. They can format a paragraph, allow you to add hidden Comments to your text, and add images within the body of your HTML text.

Note: Although this section shows hypertext procedures (HTP), all of them are also available as hypertext functions (HTF)

htp.line

Syntax

htp.line (cclear, csrc, cattributes);

Purpose

Prints the HTML tag that generates a line in the HTML document. *csrc* enables you to specify a custom image as the source of the line.

Parameters

cclear	in varchar2	DEFAULT NULL
csrc	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<HR CLEAR="cclear" SRC="csrc" cattributes>
```

htp.hr

Alias for **htp.line**

htp.nl

Syntax

htp.nl (cclear, cattributes);

Purpose

Prints the HTML tag that inserts a new line

Parameters

cclear	in varchar2	DEFAULT NULL,
cattributes	in varchar2	DEFAULT NULL,

Generates

```
<BR CLEAR="cclear" cattributes>
```

htp.br

Alias for **htp.nl**

htp.header

Syntax

htp.header (nsize, cheader, calign, cnowrap, cclear, cattributes);

Purpose

Prints the HTML tag for a heading level, with the value of the heading level assigned in the nsize parameter. Valid levels are 1 through 6.

Parameters

nsize	in integer	
cheader	in varchar2	
calign	in varchar2	DEFAULT NULL
cnowrap	in varchar2	DEFAULT NULL
cclear	in varchar2	DEFAULT NULL
cattributesi	n varchar2	DEFAULT NULL

Generates

```
<Hnsize ALIGN="calign" NOWRAP CLEAR="cclear"
   cattributes>cheader</Hnsize>
```

Example

```
htp.header (1,'Overview');
```

produces

```
<H1>Overview</H1>
```

htp.anchor

Syntax

htp.anchor (curl, ctext, cname, cattributes);

Purpose

Prints the HTML tag for an anchor to be the start or end destination of a hypertext link. This anchor can accept several attributes, but either HREF or

NAME is required. HREF specifies where to link to. NAME allows this tag to be a target of a hypertext link.

Parameters

curl	in varchar2	
ctext	in varchar2	
cname	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<A HREF="curl" NAME="cname" cattributes>ctext</A>
```

htp.anchor2

Syntax

htp.anchor2 (curl, ctext, cname, ctarget, cattributes);

Purpose

Prints the HTML tag for an anchor to be the start or end destination of a hypertext link. This anchor can accept several attributes, but either HREF or NAME is required. HREF specifies where to link to. NAME allows this tag to be a target of a hypertext link. This procedure differs from htp.anchor in that it provides a target and therefore can be used for a frame.

Parameters

curl	in varchar2	
ctext	in varchar2	
cname	in varchar2	DEFAULT NULL
ctarget	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<A HREF="curl" NAME="cname" TARGET = "ctarget" cattributes>ctext</A>
```

htp.mailto

Syntax

htp.mailto (caddress, ctext, cname, cattributes);

Purpose

Prints the HTML tag for an anchor with 'mailto' concatenated ahead of the mail address argument.

Parameters

caddress	in varchar2	
ctext	in varchar2	
cname	in varchar2	
cattributesi	n varchar2	DEFAULT NULL

Generates

```
<A HREF="mailto:caddress" cattributes>ctext</A>
```

Example

```
htp.mailto('pres@white_house.gov','Send Email to the President');
```

prints

```
<A HREF="mailto:pres@white_house.gov">Send Email to the President</A>
```

htp.img

Syntax

htp.img (curl, calign, calt, cismap, cattributes);

Purpose

Prints an HTML tag that signals the browser to load an image to be placed into the HTML page. ALT allows you to specify alternate text to be shown while the image is being loaded, or instead of the image if the browser does not support images. The ISMAP attribute indicates that the image is an image map.

Parameters

curl	in varchar2	DEFAULT NULL
calign	in varchar2	DEFAULT NULL
calt	in varchar2	DEFAULT NULL
cismap	in varchar2	DEFAULT NULL
cattributes i	n varchar2	DEFAULT NULL

Generates

```
<IMG SRC="curl" ALIGN="calign" ALT="calt" ISMAP cattributes>
```

htp.img2

Syntax

htp.img (curl, calign, calt, cismap, cusemap, cattributes);

Purpose

Prints an HTML tag that signals the browser to load an image to be placed into the HTML page. ALT allows you to specify alternate text to be shown while the image is being loaded, or instead of the image if the browser does not support images. The ISMAP attribute indicates that the image is an image map. The CUSEMAP parameter specifies a client-side image map.

Parameters

curl	in varchar2	DEFAULT NULL
calign	in varchar2	DEFAULT NULL
calt	in varchar2	DEFAULT NULL
cismap	in varchar2	DEFAULT NULL
cusemap	in varchar2	DEFAULT NULL
cattributes i	n varchar2	DEFAULT NULL

Generates

```
<IMG SRC="curl" ALIGN="calign" ALT="calt" ISMAP  USEMAP="cusemap"
cattributes>
```

htp.para

Syntax

htp.para;

Purpose

Prints an HTML tag that indicates that the text previous to it should be formatted as a paragraph.

Parameters

none

Generates

```
<P>
```

htp.paragraph

Syntax

htp.paragraph (calign, cnowrap, cclear, cattributes);

Purpose

Prints the same HTML tag as htp.para except that parameters pass in exact alignment, leading, wrapping, and attributes.

Parameters

calign	in varchar2	DEFAULT NULL
cnowrap	in varchar2	DEFAULT NULL
cclear	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<P ALIGN="calign" NOWRAP CLEAR="cclear" cattributes>
```

htp.address

Syntax

htp.address (cvalue, cnowrap, cclear, cattributes);

Purpose

Prints an HTML tag that enables you to specify address, author and signature of document

Parameters

cvalue	in varchar2	
cnowrap	in varchar2	DEFAULT NULL
cclear	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<ADDRESS CLEAR="cclear" NOWRAP cattributes>cvalue</ADDRESS>
```

htp.Comment

Syntax

htp.Comment (ctext);

Purpose

Prints an HTML tag that allows you to store Comments or lines in HTML pages. These Comments are not visible to the end user.

Parameters

ctext in varchar2

Generates

```
<!-- ctext -->
```

htp.preOpen

Syntax

htp.preOpen (cclear, cwidth, cattributes);

Purpose

Prints an HTML tag that indicates the beginning of preformatted text in the body of the HTML page.

Parameters

cclear	in varchar2	DEFAULT NULL
cwidth	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<PRE CLEAR="cclear" WIDTH="cwidth" cattributes>
```

htp.preClose

Syntax

htp.preClose;

Purpose

Prints an HTML tag that ends the preformatted section of text.

Parameters

none

Generates

```
</PRE>
```

htp.blockquoteOpen

Syntax

htp.blockquoteOpen (cnowrap, cclear, cattributes);

Purpose

Prints an HTML tag that precedes a paragraph of quoted text.

Parameters

cnowrap	in varchar2	DEFAULT NULL
cclear	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL;

Generates

```
<BLOCKQUOTE CLEAR="cclear" NOWRAP cattributes>
```

htp.blockquoteClose

Syntax

htp.blockquoteClose;

Purpose

Ends the <BLOCKQUOTE> section of quoted text.

Parameters

none

Generates

```
</BLOCKQUOTE>
```

htp.base

Syntax

htp.base(ctarget, cattributes);

Purpose

Prints an HTML tag that records the URL of the document. The ctarget attribute establishes a default window name to which all links in this document will be targeted.

Parameters

ctarget	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
BASE HREF="<current URL>" TARGET="target" cattributes>
```

htp.area

Syntax

htp.area(ccoords, cshape, chref, cnohref, ctarget, cattributes);

Purpose

Prints an HTML tag to specify the shape of a client-side image map region.

Parameters

ccoords	in varchar2	
cshape	in varchar2	DEFAULT NULL
chref	in varchar2	DEFAULT NULL
cnohref	in varchar2	DEFAULT NULL
ctarget	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<AREA COORDS="ccoords" SHAPE="cshape" HREF="chref" NOHREF TARGET="ctarget"
cattributes>
```

htp.mapOpen

Syntax

htp.mapOpen(cname, cattributes);

Purpose

Prints an HTML tag to specify a set of regions in a client-side image map.

Parameters

cname	in varchar2
cattributes	in varchar2 DEFAULT NULL

Generates

```
<MAP NAME="cname" cattributes>
```

htp.mapClose

Syntax

htp.mapClose;

Purpose

Prints an HTML tag that ends the definition of a client-side image map.

Parameters

none

Generates

```
</MAP>
```

htp.bgsound

Syntax

htp.bgsound(csrc, cloop, cattributes);

Purpose

Prints an HTML tag to include background sound for a web page.

Parameters

csrc	in varchar2
cloop	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
<BGSOUND SRC="csrc" LOOP="cloop" cattributes>
```

htp.div

Syntax

htp.div(calign, cattributes);

Purpose

Prints an HTML tag to create document divisions.

Parameters

calign	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
<DIV ALIGN="calign" cattributes>
```

htp.listingOpen

Syntax

htp.listingOpen;

Purpose

Prints an HTML tag to indicate the beginning of fixed-width text in the body of an HTML page.

Parameters

none

Generates

```
<LISTING>
```

htp.listingClose

Syntax

htp.listingClose;

Purpose

Prints an HTML tag to end the fixed-width section of text.

Parameters

none

Generates

```
</LISTING>
```

htp.nobr

Syntax

htp.nobr(ctext);

Purpose

Prints an HTML tag to turn off line-breaking with a section of text.

Parameters

ctext in varchar2

Generates

```
<NOBR>ctext</NOBR>
```


ORCL000878

htp.wbr

Syntax

htp.wbr;

Purpose

Prints an HTML tag to insert a soft linebreak within a section of NOBR text.

Parameters

none

Generates

```
<WBR>
```

htp.center

Syntax

htp.center(ctext);

Purpose

Prints a pair of HTML tags to center a section of text within a web page.

Parameters

ctexy in varchar2

Generates

```
<CENTER>ctext</CENTER>
```

htp.centerOpen

Syntax

htp.centerOpen;

Purpose

Prints an HTML tag to open a centered section of text within a web page.

Parameters

none

Generates

```
<CENTER>
```


ORCL000879

htp.centerClose

Syntax

htp.centerClose;

Purpose

Prints an HTML tag to close a centered section of text within a web page.

Parameters

none

Generates

```
</CENTER>
```

htp.dfn

Syntax

htp.dfn(ctext);

Purpose

Prints a pair of HTML tags that specify the text they surround is rendered as italics.

Parameters

ctext	in varchar2

Generates

```
<DFN>ctext</DFN>
```

htp.big

Syntax

htp.big(ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround is rendered using a big font.

Parameters

ctext	in varchar2
cattributes	in varchar2 DEFAULT NULL

ORCL000880

Generates

```
<BIG cattributes>ctext</BIG>
```

htp.small

Syntax

htp.small(ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround is rendered using a small font.

Parameters

ctext	in varchar2
cattributes	in varchar2 DEFAULT NULL

Generates

```
<SMALL cattributes>ctext</SMALL>
```

htp.sub

Syntax

htp.sub(ctext, calign, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround is rendered as a subscript.

Parameters

ctext	in varchar2
calign	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
<SUB ALIGN="calign" cattributes>ctext</SUB>
```

htp.sup

Syntax

htp.sup(ctext, calign, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround is rendered using as a superscript.

Parameters

ctext	in varchar2
calign	in varchar2
cattributes	in varchar2 DEFAULT NULL

Generates

```
<SUP ALIGN="calign" cattributes>ctext</SUP>
```

htp.basefont

Syntax

htp.basefont(nsize);

Purpose

Prints an HTML tag that specifies the base font size for a web page.

Parameters

nsize	in integer

Generates

```
<BASEFONT SIZE="nsize">
```

htp.fontOpen

Syntax

htp.fontOpen(ccolor, cface, csize, cattributes);

Purpose

Prints an HTML tag that indicates the beginning of a section of text with the specified font characteristics.

Parameters

ccolor	in varchar2 DEFAULT NULL
cface	in varchar2 DEFAULT NULL
csize	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
<FONT COLOR="ccolor" FACE="cface" SIZE="csize" cattributes>
```

htp.fontClose

Syntax

htp.fontClose;

Purpose

Prints an HTML tag that indicates the end of a section of text with the specified font characteristics.

Parameters

none

Generates

```
</FONT>
```

htp.plaintext

Syntax

htp.plaintext(ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround be rendered with fixed-width type.

Parameters

ctext	in varchar2
cattributes	in varchar2 DEFAULT NULL

Generates

```
<PLAINTEXT cattributes>ctext</PLAINTEXT>
```

htp.s

Syntax

htp.s(ctext, cattributes);

ORCL000883

Purpose

Prints a pair of HTML tags that specify the text they surround be rendered in strikethrough type.

Parameters

ctext	in varchar2
cattributes	in varchar2 DEFAULT NULL

Generates

```
<S cattributes>ctext</S>
```

htp.strike

Syntax

htp.strike(ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround be rendered in strikethrough type.

Parameters

ctext	in varchar2
cattributesi	n varchar2 DEFAULT NULL

Generates

```
<STRIKE cattributes>ctext</STRIKE>
```

Frame Tags

htp.framesetOpen

Syntax

htp.framesetOpen(crows, ccols, cattributes);

Purpose

Prints an HTML tag to open a container of web page frames.

Parameters

crows	in varchar2 DEFAULT NULL
ccols	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
<FRAMESET ROWS="crows" COLS="ccols" cattributes>
```

htp.framesetClose

Syntax

htp.framesetClose;

Purpose

Prints an HTML tag to close a container of web page frames.

Parameters

none

Generates

```
</FRAMESET>
```

htp.frame

Syntax

htp.frame(csrc, cname, cmarginwidth, cmarginheight, cscrolling, cnoresize, cattributes);

Purpose

Prints an HTML tag that defines a single frame within a frameset.

Parameters

csrc	in varchar2
cname	in varchar2 DEFAULT NULL
cmarginwidth	in varchar2 DEFAULT NULL
cmarginheight	in varchar2 DEFAULT NULL
cscrolling	in varchar2 DEFAULT NULL
cnoresizei	n varchar2 DEFAULT NULL
cattributes i	n varchar2 DEFAULT NULL

Generates

```
<FRAME SRC="csrc" NAME="cname" MARGINWIDTH="cmarginwidth"
MARGINHEIGHT="cmarginheight" SCROLLING="cscrolling" NORESIZE cattributes>
```

htp.noframesOpen

Syntax

htp.noframesOpen;

Purpose

Prints an HTML tag to open a container of content which is viewable by non-Frame-capable web browsers.

Parameters

none

Generates

```
<NOFRAMES>
```

htp.noframesClose

Syntax

htp.noframesClose;

Purpose

Prints an HTML tag to close a container of content which is viewable by non-Frame-capable web browsers.

Parameters

none

Generates

```
</NOFRAMES>
```

List Tags

List tags allow you to display information in any of the following ways:

- *ordered*: these lists have numbered items
- *unordered*: these lists have bullets to mark each item

- *definition*: these lists alternate a term with its definition

Note: All the hypertext procedures (HTP) shown in this section are also available as hypertext functions (HTF).

htp.listHeader

Syntax

htp.listHeader (ctext, cattributes);

Purpose

Prints an HTML tag at the beginning of the list

Parameters

ctext	in varchar2
cattributes	in varchar2 DEFAULT NULL

Generates

```
<LH cattributes>ctext</LH>
```

htp.listItem

Syntax

htp.listItem (ctext, cclear, cdingbat, csrc, cattributes);

Purpose

Prints an HTML tag that formats a listed item.

Parameters

ctext	in varchar2 DEFAULT NULL
cclear	in varchar2 DEFAULT NULL
cdingbat	in varchar2 DEFAULT NULL
csrc	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
<LI CLEAR="cclear" DINGBAT="cdingbat" SRC="csrc" cattributes>ctext
```


ORCL000887

htp.ulistOpen

Syntax

htp.ulistOpen (cclear, cwrap, cdingbat, csrc, cattributes);

Purpose

Prints an HTML tag that is used to open an unordered list that presents listed items separated by white space and marked off by bullets.

Parameters

cclear	in varchar2	DEFAULT NULL
cwrap	in varchar2	DEFAULT NULL
cdingbat	in varchar2	DEFAULT NULL
csrc	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<UL CLEAR="cclear" WRAP="cwrap" DINGBAT="cdingbat" SRC="csrc"
cattributes>
```

htp.ulistClose

Syntax

htp.ulistClose;

Purpose

Prints an HTML tag that ends the unordered list.

Parameters

none

Generates

```
</UL>
```

htp.olistOpen

Syntax

htp.olistOpen (cclear, cwrap, cattributes);

Purpose

Prints an HTML tag that is used to open an ordered list that presents listed items marked off with numbers.

Parameters

cclear	in varchar2	DEFAULT NULL
cwrap	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<OL CLEAR="cclear" WRAP="cwrap" cattributes>
```

htp.olistClose

Syntax

htp.olistClose;

Purpose

Prints an HTML tag that ends an ordered list.

Parameters

none

Generates

```
</OL>
```

htp.dlistOpen

Syntax

htp.dlistOpen (cclear, cattributes);

Purpose

Prints an HTML tag that starts a definition list

Parameters

cclear	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<DL CLEAR="cclear" cattributes>
```

htp.dlistClose

Syntax

htp.dlistClose

Purpose

Prints an HTML tag that Ends a definition list

Parameters

none

Generates

```
</DL>
```

htp.dlistDef

Syntax

htp.dlistDef (ctext, cclear, cattributes);

Purpose

Prints an HTML tag that is used to insert terms, and their corresponding definitions in an indented list format. The htp.dlistTerm must immediately follow this tag.

Parameters

ctext	in varchar2	DEFAULT NULL
clear	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<DD CLEAR="cclear" cattributes>ctext
```

htp.dlistTerm

Syntax

htp.dlistTerm (ctext, cclear, cattributes);

Purpose

Prints an HTML tag used to insert the definition term inside the definition list. This tag must immediately follow the htp.dlistDef.

Parameters

ctext	in varchar2	DEFAULT NULL
cclear	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<DT CLEAR="cclear" cattributes>ctext
```

htp.menulistOpen

Syntax

htp.menulistOpen;

Purpose

Prints an HTML tag that begins a list that presents one line per item, and appears more compact than an unordered list. The htp.listItem will follow this tag.

Parameters

none

Generates

```
<MENU>
```

htp.menulistClose

Syntax

htp.menulistClose;

Purpose

Prints an HTML tag that ends a menu list.

Paramenters

none

Generates

```
</MENU>
```

htp.dirlistOpen

Syntax

htp.dirlistOpen;

Purpose

Prints an HTML tag that begins a directory list. This presents information in a list of items that contain up to 20 characters. Items in this list are typically arranged in columns, typically 24 characters wide. The <LI> or htp.listItem must appear directly after you use this tag.

Parameters

none

Generates

```
<DIR>
```

htp.dirlistClose

Syntax

htp.dirlistClose;

Purpose

Prints an HTML tag that closes the directory list tag, htp.dirlistOpen.

Parameters

none

Generates

```
</DIR>
```

Character Format Tags

The character format tags are used to specify or alter the appearance of the marked text. Character format tags have opening and closing elements, and affect only the text that they surround.

Character format tags give hints to the browser as to how a character or character string should appear, but each browser determines its actual appearance. Essentially, they place text into categories such that all text in a given category is given the same special treatment, but the browser determines what that

treatment is. For example, the HTML string `<STRONG>Here is some text</STRONG>` might appear as bold in some browsers, or might flash instead.

If a specific text attribute, such as bold is desired, a physical format tag may be necessary. See the section, "Physical Format Tags," for more information.

Note: All the hypertext procedures (HTP) shown in this section are also available as hypertext functions (HTF).

htp.cite

Syntax

htp.cite (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround as a citation. Usually rendered as italics.

Parameters

ctext	in varchar2
cattributes	in varchar2DEFAULT NULL

Generates

```
<CITE cattributes>ctext</CITE>
```

htp.code

Syntax

htp.code (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround as an example of code output. Usually rendered in monospace format, e.g. Courier.

Parameters

ctext	in varchar2
cattributes	in varchar2DEFAULT NULL

Generates

```
<CODE cattributes>ctext</CODE>
```

htp.emphasis

Syntax

htp.emphasis (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround as requiring typographic emphasis. This tag is equivalent to htp.em. Usually rendered as italics.

Parameters

ctext	in varchar2
cattributes	in varchar2DEFAULT NULL

Generates

```
<EM cattributes>ctext</EM>
```

htp.em

Alias for htp.emphasis

htp.keyboard

Syntax

htp.keyboard (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround as text typed in by the user, which usually is rendered as monospace. This tag is equivalent to htp.kbd.

Parameters

ctext	in varchar2
cattributes	in varchar2DEFAULT NULL

Generates

```
<KBD cattributes>ctext</KBD>
```

htp.kbd

Alias for htp.keyboard

ORCL000894

htp.sample

Syntax

htp.sample (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround as a sequence of literal characters that must be typed in the exact sequence in which they appear. Usually rendered as monospace font.

Parameters

ctext	in varchar2
cattributesi	n varchar2DEFAULT NULL

Generates

```
<SAMP cattributes>ctext</SAMP>
```

htp.strong

Syntax

htp.strong (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround as needing strong typographic emphasis. Usually rendered as bold.

Parameters

ctext	in varchar2
cattributes	in varchar2DEFAULT NULL

Generates

```
<STRONG cattributes>ctext</STRONG>
```

htp.variable

Syntax

htp.variable (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround as a variable name, or a variable that might be entered by the user. Usually rendered as italics.

Parameters

ctext	in varchar2	
cattributes	in varchar2	DEFAULT NULL

Generates

```
<VAR cattributes>ctext</VAR>
```

Physical Format Tags

The physical format tags are used to specify the format of the marked text.

Note: All the hypertext procedures (HTP) shown in this section are also available as hypertext functions (HTF).

htp.bold

Syntax

htp.bold (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround is to be rendered as boldface.

Parameters

ctext	in varchar2	
cattributes	in varchar2	DEFAULT NULL

Generates

```
<B cattributes>ctext</B>
```

htp.italic

Syntax

htp.italic (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround is to be rendered as italics.

Parameters

ctext	in varchar2	
cattributesi	n varchar2	DEFAULT NULL

Generates

```
<I cattributes>ctext</I>
```

htp.teletype

Syntax

htp.teletype (ctext, cattributes);

Purpose

Prints a pair of HTML tags that specify the text they surround is to be rendered in a fixed width typewriter font, e.g. Courier.

Parameters

ctext	in varchar2	
cattributesi	n varchar2	DEFAULT NULL

Generates

```
<TT cattributes>ctext</TT>
```

Form Tags

The form tags are used to create and manipulate an HTML form. Forms are used to allow interactive data exchange between a Web Browser and a CGI program.

Forms can have the following types of elements:

- Input: used for a large variety of types of input fields, for example:
  - single line text
  - single line password fields
  - checkboxes
  - radio buttons
  - submit buttons
- Text area: used to create a multi-line input field.

- Select: used to allow the user to chose one or more of a set of alternatives described by textual labels. Usually rendered as a pulldown, pop up, or a fixed size list.

Note: All the hypertext procedures (HTP) shown in this section are also available as hypertext functions (HTF).

htp.formOpen

Syntax

htp.formOpen (curl, cmethod, ctarget, cenctype, cattributes);

Purpose

Prints an HTML tag that starts the form. The *curl* value is required and is the URL of the CGI script, normally owa, to which the contents of the Form will be sent. The method is either "GET" or "POST."

Parameters

curl	in varchar2
cmethod	in varchar2DEFAULT 'POST
ctarget	in varchar2
cenctype	in varchar2DEFAULT NULL
cattributes	in varchar2DEFAULT NULL

Generates

```
<FORM ACTION="curl" METHOD="cmethod" TARGET="ctarget" ENCTYPE="cenctype"
cattributes>
```

htp.formClose

Syntax

htp.formClose;

Purpose

Prints an HTML tag that closes the <FORM> tag

Parameters

none

Generates

```
</FORM>
```

htp.formCheckbox

Syntax

htp.formCheckbox (cname, cvalue, cchecked, cattributes);

Purpose

Prints an HTML tag that inserts a checkbox which the user can toggle off or on.

Parameters

cname	in varchar2	
cvalue	in varchar2	DEFAULT 'on'
cchecked	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<INPUT TYPE="checkbox" NAME="cname" VALUE="cvalue" CHECKED cattributes>
```

htp.formHidden

Syntax

htp.formHidden (cname, cvalue, cattributes);

Purpose

Prints an HTML tag that sends the content of a field along with a submitted form. The field is not visible to the end user.

Parameters

cname	in varchar2	
cvalue	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<INPUT TYPE="hidden" NAME="cname" VALUE="cvalue" cattributes>
```

htp.formImage

Syntax

htp.formImage (cname, csrc, calign, cattributes);

ORCL000899

Purpose

Prints an HTML tag that creates an image field that can be clicked on causing the Form to be immediately submitted. The coordinates of the selected point are measured in pixels, and returned (along with other contents of the form) in two name/value pairs. The x-coordinate is submitted under the name of the field with ".x" appended, and the y-coordinate with the ".y" appended. Any value attribute is ignored. The image itself is specified by the CSRC attribute.

Parameters

cname	in varchar2	
csrc	in varchar2	
calign	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL;

Generates

```
<INPUT TYPE="image" NAME="cname" SRC="csrc" ALIGN="calign" cattributes>
```

htp.formPassword

Syntax

htp.formPassword (cname, csize, cmaxlength, cvalue, cattributes);

Purpose

Prints an HTML tag that creates a single line text entry field. Text will not be displayed as it is entered. When the user enters a password, characters are represented by asterisks on single line text entry field.

Parameters

cname	in varchar2	
csize	in varchar2	
cmaxlength	in varchar2	DEFAULT NULL
cvalue	in varchar2	DEFAULT NULL
cattributes i	n varchar2	DEFAULT NULL

Generates

```
<INPUT TYPE="password" NAME="cname" SIZE="csize" MAXLENGTH="cmaxlength"
VALUE="cvalue" cattributes>
```

htp.formRadio

Syntax

htp.formRadio (cname, cvalue, cchecked, cattributes);

Purpose

Prints an HTML tag that inserts a radio button on the HTML Form. Used to create a set of radio buttons, each representing a different value, only one of which will be toggled on by the user. Each radio button field should have the same name. Only the selected radio button will generate a name/value pair in submitted data area. This will require an explicit VALUE attribute.

Parameters

cname	in varchar2	
cvalue	in varchar2	
cchecked	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<INPUT TYPE="radio" NAME="cname" VALUE="cvalue" CHECKED cattributes>
```

htp.formReset

Syntax

htp.formReset (cvalue, cattributes);

Purpose

Prints an HTML tag that creates a RESET button that, when selected, resets all the form fields to their initial values.

Parameters

cvalue	in varchar2	DEFAULT 'Reset'
cattributes	in varchar2	DEFAULT NULL

Generates

```
<INPUT TYPE="reset" VALUE="cvalue" cattributes>
```

htp.formSubmit

Syntax

htp.formSubmit (cname, cvalue, cattributes);

Purpose

Prints an HTML tag that creates a button that, when selected, submits the form. If a SUBMIT button is selected to submit the Form, and that button has a name attribute specified, the submit button then contributes a name/value pair to the submitted data.

Parameters

cname in varchar2DEFAULT NULL
cvalue in varchar2DEFAULT 'Submit'
cattributes in varchar2DEFAULT NULL

Generates

```
<INPUT TYPE="submit" NAME="cname" VALUE="cvalue" cattributes>
```

htp.formText

Syntax

htp.formText (cname, csize, cmaxlength, cvalue, cattributes);

Purpose

Prints an HTML tag that creates a field for a single line of text.

Parameters

cname	in varchar2	
csize	in varchar2	DEFAULT NULL
cmaxlength	in varchar2	DEFAULT NULL
cvalue	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<INPUT TYPE="text" NAME="cname" SIZE="csize" MAXLENGTH="cmaxlength"
VALUE="cvalue" cattributes>
```

htp.formSelectOpen

Syntax

htp.formSelectOpen (cname, cprompt, nsize, cattributes);

Purpose

Prints an HTML tag that begins a Select list of alternatives. Contains the attribute NAME which specifies the name that will be submitted as a name/value pair.

Parameters

cname	in varchar2
cprompt	in varchar2 DEFAULT NULL
nsize	in integer DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
cprompt <SELECT NAME="cname" PROMPT="cprompt" SIZE="nsize" cattributes>
```

Example

```
htp.formSelectOpen('greatest_player';
'Pick the greatest player:');
htp.formSelectOption('Messier');
htp.formSelectOption('Howe');
htp.formSelectOption('Hull');.
htp.formSelectOption('Gretzky');.
htp.formSelectClose;

generates:

Pick the greatest player:
 <SELECT NAME="greatest_player">
 <OPTION>Messier
 <OPTION>Howe
 <OPTION>Hull
 <OPTION>Gretzky
</SELECT>
```

Note: See htp.formSelectOption and htp.formSelectClose.

htp.formSelectOption

Syntax

htp.formSelectOption (cvalue, cselected, cattributes);

Purpose

Prints an HTML tag that represents one choice in the Select element.

Parameters

cvalue	in varchar2
cselected	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
<OPTION SELECTED cattributes>cvalue
```

See example for htp.formSelectOpen.

Example

See htp.formSelectOpen.

htp.formSelectClose

Syntax

htp.formSelectClose;

Purpose

Prints an HTML tag that ends a Select list of alternatives.

Parameters

none

Generates

```
</SELECT>
```

Example

See htp.formSelectOpen.

htp.formTextarea

Syntax

htp.formTextarea (cname, nrows, ncolumns, calign, cattributes);

Purpose

Prints an HTML tag that creates a text field that has no predefined text in the text area. Used to enable the user to enter several lines of text.

Parameters

cname	in varchar2	
nrows	in integer	
ncolumns	in integer	
calign	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign"
cattributes></TEXTAREA>
```

htp.formTextarea2

Syntax

htp.formTextarea (cname, nrows, ncolumns, calign, cwrap, cattributes);

Purpose

Prints an HTML tag that creates a text field that has no predefined text in the text area. Used to enable the user to enter several lines of text. The difference between this and formTextarea is the cwrap parameter, which specifies a wrap style.

Parameters

cname	in varchar2
nrows	in integer
ncolumns	in integer
calign	in varchar2 DEFAULT NULL
cwrap	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign"
WRAP="cwrap" cattributes></TEXTAREA>
```

htp.formTextareaOpen

Syntax

htp.formTextareaOpen (cname, nrows, ncolumns, calign, cattributes);

Purpose

Prints an HTML tag that opens a text area where you can insert predefined text that will always appear in the text field.

Parameters

cname	in varchar2
nrows	in integer
ncolumns	in integer
calign	in varchar2 DEFAULT NULL
cattributes	in varchar2 DEFAULT NULL

Generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign"
cattributes>
```

htp.formTextareaOpen2

Syntax

htp.formTextareaOpen (cname, nrows, ncolumns, calign, cwrap, cattributes);

Purpose

Prints an HTML tag that opens a text area where you can insert predefined text that will always appear in the text field. The difference between this and formTextareaOpen is the cwrap parameter, which specifies a wrap style.

Parameters

cname	in varchar2	
nrows	in integer	
ncolumns	in integer	
calign	in varchar2	DEFAULT NULL
cwrap	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign" WRAP =
"cwrap" cattributes>
```

htp.formTextareaClose

Syntax

htp.formTextareaClose;

Purpose

Prints an HTML tag that ends TextArea field

Parameters

none

Generates

```
</TEXTAREA>
```

Table Tags

The Table tags allow the user to insert tables and manipulate the size and columns of the table in a document.

Note: All the hypertext procedures (HTP) shown in this section are also available as hypertext functions (HTF).

htp.tableOpen

Syntax

htp.tableOpen (cborder, calign, cnowrap, cclear, cattributes);

Purpose

Prints an HTML tag that begins an HTML table.

Parameters

cborder	in varchar2	DEFAULT NULL
calign	in varchar2	DEFAULT NULL
cnowrap	in varchar2	DEFAULT NULL
cclear	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL;

Generates

```
<TABLE "cborder" NOWRAP ALIGN="calign" CLEAR="cclear" cattributes>
```

htp.tableClose

Syntax

htp.tableClose;

Purpose

Prints an HTML tag that ends an HTML table.

Parameters

none

Generates

```
</TABLE>
```

htp.tableCaption

Syntax

htp.tableCaption (ccaption, calign, cattributes);

Purpose

Prints an HTML tag that places a caption in the inserted table.

Parameters

ccaption	in varchar2	
calign	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<CAPTION ALIGN="calign" cattributes>ccaption</CAPTION>
```

htp.tableRowOpen

Syntax

htp.tableRowOpen (calign, cvalign,cdp, cnowrap, cattributes);

Purpose

Prints an HTML tag that inserts a row tag into a table.

Parameters

calign	in varchar2	DEFAULT NULL
cvalign	in varchar2	DEFAULT NULL
cdp	in varchar2	DEFAULT NULL
cnowrap	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<TR ALIGN="calign" VALIGN="cvalign" DP="cdp" NOWRAP catttributes>
```

htp.tableRowClose

Syntax

htp.tableRowClose;

Purpose

Prints an HTML tag that ends a row in a table.

Parameters

none

Generates

```
</TR>
```

htp.tableHeader

Syntax

htp.tableHeader (cvalue, calign, cdp, cnowrap, crowspan, ccolspan, cattributes);

Purpose

Prints an HTML tag that inserts a table header.

Parameters

cvalue	in varchar2	DEFAULT NULL
calign	in varchar2	DEFAULT NULL
cdp	in varchar2	DEFAULT NULL
cnowrap	in varchar2	DEFAULT NULL
crowspan	in varchar2	DEFAULT NULL
ccolspan	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<TH ALIGN="calign" DP="cdp" ROWSPAN="crowspan" COLSPAN="ccolspan" NOWRAP
cattributes>cvalue</TH>
```

htp.tableData

Syntax

htp.tableData (cvalue, calign, cdp, crowspan, ccolspan, cnowrap, cattributes);

Purpose

Prints an HTML tag that inserts data into the rows and columns of a selected table.

Parameters

cvalue	in varchar2	DEFAULT NULL
calign	in varchar2	DEFAULT NULL
cdp	in varchar2	DEFAULT NULL
crowspan	in varchar2	DEFAULT NULL
ccolspan	in varchar2	DEFAULT NULL

cnowrap	in varchar2	DEFAULT NULL
cattributes	in varchar2	DEFAULT NULL

Generates

```
<TD ALIGN="calign" DP="cdp" ROWSPAN="crowspan" COLSPAN="ccolspan" NOWRAP
cattributes>cvalue</TD>
```

OWA_UTIL Package

The OWA_UTIL Package is a set of useful utility procedures built on top of hypertext functions and hypertext procedures.

owa_util.signature

Syntax

owa_util.signature;

Purpose

Prints an HTML line followed by a signature line on the HTML document.

Parameters

none

Generates

Prints a signature line in the HTML document that might look like the following: "This page was produced by the **Oracle PL/SQL Agent** on August 9, 1995 09:30"

owa_util.signature (cname)

Syntax

owa_util.signature (cname);

Purpose

Allows the programmer to create a signature line on the bottom of the HTML document that has a hypertext link to view the PL/SQL source for that procedure.

Parameters

cname in varchar2

Generates

Prints a signature line on the bottom of the HTML document that has a link to the actual PL/SQL source for that procedure. The link calls the procedure `showsource`. The line would look like the following:

"This page was produced by the **Oracle PL/SQL Agent** on 6/14/95 09:30"
View PL/SQL Source

owa_util.showsource (cname)

Syntax

owa_util.showsource (cname);

Purpose

Prints the source of the specified PL/SQL procedure, function, or package. If a procedure or function which belongs to a package is specified, then the entire package is displayed.

Parameters

cname in varchar2

Generates

Generates the source code of the specified PL/SQL procedure.

owa_util.showpage

Syntax

owa_util.showpage;

Purpose

This procedure allows a user to view the HTML output of a PL/SQL procedure call from SQL*Plus, SQL*DBA, or Oracle Server Manager. The PL/SQL procedure must use HTP and/or HTF to generate the HTML page, and owa_util.showpage must be issued after the PL/SQL procedure has been called and before any other HTP or HTF subprograms are directly or indirectly called. This method is useful for generating pages filled with static data.Purpose

Note that this procedure uses dbms_output and thus is limited to 255 characters per line and an overall buffer size of 1,000,000 bytes.

Parameters

none

Generates

One can use this procedure to generate static pages in SQL*Plus that can then be accessed as a standard HTML page. For example:

```
SQL>set serveroutput on
SQL>spool gretzky.html
SQL>execute hockey.pass('Gretzky")
SQL>execute owa_util.showpage
SQL>exit
```

This would generate an HTML page which could be accessed from Web clients.

owa_util.get_cgi_env(function)

Syntax

owa_util.get_cgi_env(param_name in varchar2);

Purpose

Allows programmer to retrieve the value of the specified CGI environment variable in the PL/SQL procedure.

Note that param_name is case-insensitive.

Parameters

param_name in varchar2

Generates

Returns value of specified CGI environment variable for PL/SQL procedure. If the value is not set, returns *null*

owa_util.print_cgi_env

Syntax

owa_util.print_cgi_env;

Purpose

Enables programmer to print all of the CGI environment variables made available by the PL/SQL Agent to the PL/SQL procedures. This utility is good for testing purposes.

Parameters

none

Generates

Prints CGI environment variables made available by the PL/SQL Agent to the PL/SQL procedures.

owa_util.mime_header

Syntax

owa_util.mime_header(ccontent_type, bclose_header);

Purpose

Enables programmer to change the default MIME header that the PL/SQL Agent returns. This *must* come before any htp.print or htp.prn calls in order to signal the PL/SQL Agent not to use the default. If bcloseheader is TRUE, two newlines are sent, which closes the HTTP header. Otherwise, one newline is sent, and the HTTP header is still open.

Parameters

ccontent_type	in varchar2
bclose_header	in boolean DEFAULT TRUE

Generates

```
Content-type: <ccontent_type>\n\n
```

owa_util.redirect_url

Syntax

owa_util.redirect_url(curl, bclose_header);

Purpose

Enables programmer to specify that the WebServer visit a specified URL. The URL may specify either a Web page that is returned or a program that is executed. This *must* come before any htp.print or htp.prn calls in order to signal the PL/SQL Agent to do the redirect. If bcloseheader is TRUE, two newlines are sent, which closes the HTTP header. Otherwise, one newline is sent, and the HTTP header is still open.

Parameters

curl	in varchar2
bclose_header	in boolean DEFAULT TRUE

Generates

```
Location: <curl>\n\n
```

owa_util.status_line

Syntax

owa_util.status_line(nstatus, creason, bclose_header);

Purpose

Enables programmer to send a standard HTTP status code to the client. This *must* come before any htp.print or htp.prn calls, so that the status code is returned as part of the header, rather than as "content data." If bcloseheader is TRUE, two newlines are sent, which closes the HTTP header. Otherwise, one newline is sent, and the HTTP header is still open.

Parameters

nstatus	in integer,
creason	in varchar2 DEFAULT NULL
bclose_header	in boolean DEFAULT TRUE

Generates

```
Status: <nstatus> <creason>\n\n
```

owa_util.http_header_close

Syntax

owa_util.http_header_close;

Purpose

Outputs a newline to close the HTTP header. Use this procedure if you have not explicitly closed the header before by specifying or defaulting bclose_header in a previous header component. The header *must* be closed before any htp.print or htp.prn calls.

Parameters

none

Generates

A newline, which terminates the header.

owa_util.get_owa_service_path (Function)

Syntax

owa_util.get_owa_service_path;

Purpose

Returns the name of the currently active path with its full virtual path, plus the currently active DCD. For example, a call to get_owa_service_path could return `/ows-bin/myservice/owa/`.

Parameters

none

Returns

The DCD path. The datatype is varchar2.

owa_util.tableprint

Syntax

owa_util.tablePrint;

Purpose

Enables programmers to print Oracle tables as either preformatted or HTML tables, depending upon Web browser capabilities. Note that RAW COLUMNS are supported, however LONG RAW are not. References to LONG RAW columns will print the result 'Not Printable'. In this case, `cattributes` is the second, rather than the last, parameter.

Parameters

ctable in varchar2
cattributes in varchar2 DEFAULT NULL
ntable_type in integer DEFAULT HTML_TABLE
ccolumns in varchar2 DEFAULT '*'
cclauses in varchar2 DEFAULT NULL
ccol_aliasesin varchar2 DEFAULT NULL
nrow_min in number DEFAULT 0
nrow_max in number DEFAULT NULL

Parameters

Note that ntable_type can be either owa_util.html_table or owa_util.pre_table.

Generates

Prints out either a preformatted or HTML table.

Returns

True or False as to whether there are more rows available beyond the nrow_max requested.

Example

For browsers that don't support HTML tables, create the following procedure:

```
create or replace procedure showemps is
  ignore_more boolean;
begin
  ignore_more := owa_util.tablePrint('emp', 'BORDER', OWA_UTIL.PRE_TABLE);
end;
 and requesting a URL like this example: http://myhost:8080/ows-bin/hr/owa/
showemps

returns to the client:
<PRE>
```

EMPNO	ENAME	JOB	MGR	HIREDATE	SAL	COMM	DEPTNO
7369	SMITH	CLERK	7902	17-DEC-80	800	.	20
7499	ALLEN	SALESMAN	7698	20-FEB-81	1600	300	30
7521	WARD	SALESMAN	7698	22-FEB-81	1250	500	30
7566	JONES	MANAGER	7839	02-APR-81	2975		20
7654	MARTIN	SALESMAN	7698	28-SEP-81	1250	1400	30
7698	BLAKE	MANAGER	7839	01-MAY-81	2850		30
7782	CLARK	MANAGER	7839	09-JUN-81	2450		10
7788	SCOTT	ANALYST	7566	09-DEC-82	3000		20
7839	KING	PRESIDENT		17-NOV-81	5000		10
7844	TURNER	SALESMAN	7698	08-SEP-81	1500	0	30
7876	ADAMS	CLERK	7788	12-JAN-83	1100		20
7900	JAMES	CLERK	7698	03-DEC-81	950		30
7902	FORD	ANALYST	7566	03-DEC-81	3000		20
7934	MILLER	CLERK	7782	23-JAN-82	1300		10

```
</PRE>
```

To view just the employees in department 10, and only their employee ids, names, and salaries, create the following procedure:

```
create or replace procedure showemps_10 is
   ignore_more boolean;
begin
 ignore_more := owa_util.tablePrint
('EMP', 'BORDER', OWA_UTIL.PRE_TABLE,
       'empno, ename, sal',
       'where deptno=10 order by empno',
       'Employee Number, Name, Salary');
end;
```

A request for a URL like `http://myhost:8080/ows-bin/hr/owa/showemps_10` would return the following to the client:

```
<PRE>
-------------------------------------
```

```
| Employee Number |  Name  | Salary |
-------------------------------------
| 7782            | CLARK  | 2450   |
| 7839            | KING   | 5000   |
| 7934            | MILLER | 1300   |
-------------------------------------
</PRE>
```

For browsers that do support HTML tables, to view the department table in an HTML table, create the following procedure:

```
create or replace procedure showdept is
 ignore_more boolean;
begin
   ignore_more := owa_util.tablePrint('dept', 'BORDER');
end;
```

A request for a URL like `http://myhost:8080/ows-bin/hr/owa/showdept` would return the following to the client:

```
<TABLE BORDER>
<TR>
<TH>DEPTNO</TH>
<TH>DNAME</TH>
<TH>LOC</TH>
</TR>
<TR>
<TD ALIGN="LEFT">10</TD>
<TD ALIGN="LEFT">ACCOUNTING</TD>
<TD ALIGN="LEFT">NEW YORK</TD>
</TR>
<TR>
<TD ALIGN="LEFT">20</TD>
<TD ALIGN="LEFT">RESEARCH</TD>
<TD ALIGN="LEFT">DALLAS</TD>
</TR>
<TR>
<TD ALIGN="LEFT">30</TD>
<TD ALIGN="LEFT">SALES</TD>
<TD ALIGN="LEFT">CHICAGO</TD>
</TR>
<TR>
<TD ALIGN="LEFT">40</TD>
<TD ALIGN="LEFT">OPERATIONS</TD>
<TD ALIGN="LEFT">BOSTON</TD>
</TR>
</TABLE>
```

which a Web browser can format to look like this:

DEPTNO	DNAME	LOC
10	ACCOUNTING	NEW YORK
20	RESEARCH	DALLAS
30	SALES	CHICAGO

Customized Extensions to HTP Packages

The design of the hypertext procedure and function packages allows you to use customized extensions. Therefore, as the HTML standard changes, you can add new functionality similar to the hypertext procedure and function packages to reflect those changes.

Here is an example of customized packages using non-standard <BLINK> and imaginary <SHOUT>tags:

```
create package nsf as
function blink(cbuf in varchar2) return varchar2;
function shout(cbuf in varchar2) return varchar2;
end;
create package body nsf as
function blink(cbuf in varchar2) return varchar2 is
begin return ('<BLINK>' || cbuf || '</BLINK>');
end;
function shout(cbuf in varchar2) return varchar2 is
begin return ('<SHOUT>' || cbuf || '</SHOUT>');
end;
end;

create package nsp as
procedure blink(cbuf  in varchar2);
procedure shout(cbuf  in varchar2);
end;
create package body nsp as
procedure blink(cbuf  in varchar2) is
begin htp.print(nsf.blink(cbuf)); end;
procedure shout(cbuf  in varchar2) is
begin htp.print(nsf.shout(cbuf)); end;
end;
```

Now you can begin to use these procedures and functions in your own procedure.

```
create procedure nonstandard as
begin
nsp.blink('Gee this hurts my eyes!');
htp.print('And I might ' || nsf.shout('get mad!'));
end;
```

For more examples of using the PI/SQL Web Toolkit, see *Passing Parameters to PL/SQL*.

OWA_PATTERN Package

The OWA_PATTERN package enables you to do sophisticated string manipulation using regular expressions. OWA_PATTERN provides the following three operations:

- MATCH. This determines whether a regular expression exists in a string. This is a function that returns TRUE or FALSE.

- AMATCH. This is a more sophisticated variation on MATCH that lets you specify *where* in the string the match has to occur. This is a function that returns as an integer the end of the location in the string where the regular expression was found. If the regular expression is not found, it returns 0.
- CHANGE. This lets you replace the portion of the string that matched the regular expression with a new string. CHANGE can be either a procedure or a function. If a function, it returns the number of times the regular expression was found and replaced.

There are also operations that these operations use, but that you can also use directly. These are explained shortly.

The OWA_PATTERN operations all use the following three parameters:

- *line*. This is the target to be examined for a match. Despite the name, it can be more than one line of text or can be a PL/SQL table(see *PL/SQL Tables*) of type multi_line
- *pat*. This is the regular expression the functions attempts to locate in *line*. This regular expression uses the special tokens explained shortly. Note: in CHANGE, this is called *from_str*.
- *flags*. These are arguments that control how the search is to be performed.

Some of the operations take additional parameters as well.

Regular Expressions

You specify a regular expression by creating the string you want to match interspersed with various wildcard tokens and quantifiers. The wildcard tokens all match something other than themselves, and the quantifiers modify the meaning of tokens or of literals by specifying such things as how often each is to be applied.

Tokens

The wildcard tokens that are supported are as follows:

^ Matches newline or the beginning of the target.

$ Matches newline or the end of the target.

\n Matches newline.

. Matches any character except newline.

\t Matches tab.

\d Matches digits [0-9]

\D Matches non-digits [not 0-9]

\w Matches word characters (alphanumeric) [0-9, a-z, A-Z or _]

\W Matches non-word characters [not 0-9, a-z, A-Z or _]

\s Matches whitespace characters [blank, tab, or newline]

\S Matches non-whitespace characters [not blank, tab, or newline]

\b Matches "word" boundaries (between \w and \W)

\x<HEX> Matches the value in the current character set of the two hexidecimal digits.

\<OCT> Matches the value in the current character set of the two or three octal digits.

\ Followed by any character not covered by another case matches that character.

& Applies only to CHANGE. This causes the string that matched the regular expression to be included in the string that replaces it. This differs from the other tokens in that it specifies how a target is changed rather than how it is matched. This is explained further under CHANGE.

Quantifiers

Any of the above tokens except & can have its meaning extended by any of the following quantifiers. You can also apply these quantifiers to literals.

? 0 or 1 occurrence(s)

* 0 or more occurrences

+ 1 or more occurrence(s)

{n} Exactly *n* occurrences

{n,} At least *n* occurrences

{n,m} At least *n*, but not more than *m*, occurrences

Flags

In addition to targets and regular expressions, the OWA_PATTERN functions and procedures can use flags to affect how they are interpreted. The recognized flags are as follows:

i This indicates a case-insensitive search.

g This applies only to CHANGE. It indicates a global replace. That is to say, all portions of the target that match the regular expression are replaced.

Datatypes

The following special datatype is used by OWA_PATTERN.

pattern

A PL/SQL table (see *PL/SQL Tables*) of 4 byte VARCHAR2 strings, indexed by BINARY INTEGER. This is an alternative way to store your regular expression than in simple VARCHAR2 strings. The advantages of this is that you can use a pattern as both an input and output parameter. Thus, you can pass the same regular expression to several subsequent OWA_PATTERN function calls, and it only has to be parsed once.

Note: The following datatypes are used by OWA_PATTERN, but are part of the OWA_TEXT package. For information on these, see the section on OWA_TEXT: *owa_text.vc_array, owa_text.multi_line, owa_text.int_array, owa_text.row_list.*

Using MATCH, AMATCH, and CHANGE

Here is an example of MATCH.

```
MATCH ('BATMAN', 'Bat.*', i);
```

This is how the function is interpreted: BATMAN is the target where we are searching for the regular expression. Bat.* is the regular expression we are attempting to find. The period (.) indicates any character other than newline, and the asterisk (*) indicates any 0 or more of such. Therefore, this regular expression specifies that a matching target consists of 'Bat', followed by any set of characters neither ending in nor including a newline (which does not match the period). The *i* at the end is a flag indicating that case is to be ignored in the search.

This would return TRUE, indicating that a match had been found.

Note that, if multiple overlapping strings can match the regular expression, OWA_PATTERN takes the longest match.

Summaries of OWA_PATTERN Functions

MATCH, AMATCH, and CHANGE are overloaded. That is to say, there are several versions of each, distinguished by the parameters they take. Specifically, there are six versions of MATCH, and four each of AMATCH and CHANGE. This section provides a summary of all versions; the following section provides a reference on each version.

MATCH is a function that returns TRUE or FALSE depending on whether a match was found. Here is a summary of the versions of MATCH:

- The target can be either a simple VARCHAR2 string of less than 32K or a multi_line. A multi-line is described under "OWA_TEXT Datatypes". You can create a multi_line from a long string using the *stream2multi* function described under OWA_TEXT. If a multi_line is used, there is a parameter called *rlist*, after the regular expression, but before the flags. This is a list of the chunks where matches were found. Use of a VARCHAR2 implies use of one of the first four versions of this function. Use of a multi_line implies use of version 5 or 6.
- The regular expression can be either a VARCHAR2 string or a *pattern*. You can create a *pattern* from a string using the *getpat* function described later in this section. If you use a pattern, you are using one of the even-number versions of the function.
- If the line is a string and not a multi_line, then you can add an optional output parameter called *backrefs*. This goes after the regular expression, but before the flags. You cannot use *backrefs* if you pass a multi_line to MATCH, because this is the same place in the parameter list that the *rlist* parameter would go. The *backrefs* parameter is a row_list that holds each string in the target that was matched by a sequence of tokens in the regular expression. If you use *backrefs*, you are using the third or fourth version of the function.
- Given the above, you can determine which version you are using as follows: if you are using a multi_line, it is version 5 or 6. If you are using backrefs, it is version 3 or 4. Otherwise, it is version 1 or 2. Once you have determined the pair, you determine the actual version by whether the regular expression is a VARCHAR2 string (odd-numbered) or a pattern (even-numbered).

AMATCH is a function giving a number that indicates the number of characters from the beginning of the target to the end of the first match found (AMATCH stops searching after the first match). If no match is found, it returns 0. Here is a summary of the versions of AMATCH:

- Following the target, but preceding the regular expression, is the input parameter *from_loc*. This indicates how many characters from the beginning of the target the search should commence.
- The regular expression can be either a VARCHAR2 string or a pattern. You can create a pattern from a string using the *getpat* function described later in this section.. If you use a pattern, you are using one of the even-number versions of the function.

- After the regular expression, but before the flags, you can add an optional output parameter called "backrefs". This is a PL/SQL table (see *PL/SQL Tables*) that will hold each string in the target that was matched by a sequence of tokens in the regular expression. If you use *backrefs*, you are using the third or fourth version of the function.
- Given the above, you can determine which version you are using as follows: If you are using backrefs, it is version 3 or 4. Otherwise, it is version 1 or 2. Once you have determined the pair, you determine the actual version by whether the regular expression is a VARCHAR2 string (odd-numbered) or a pattern (even-numbered).

CHANGE can be either a procedure or a function, depending on how it is invoked. If a function, it returns the number of changes made. If the flag 'g' is not used, this number can be only 0 or 1. Here is a summary of the versions of CHANGE

- The target can be either a simple VARCHAR2 string of less than 32K or a multi_line.
- Unlike MATCH and AMATCH, the regular expression can only be a VARCHAR2 string, not a pattern.
- Following the regular expression is the string that is to replace it. This string can use the token ampersand (&), which indicates that the portion of the target that matched the regular expression is to be included in the expression that replaces it. For example:

```
CHANGE('Cats in pajamas','C.+in', '& red ')
```

The regular expression matches the substring 'Cats in'. It then replaces this string with '& red'. & indicates 'Cats in', since that's what matched the regular expression. Thus, this procedure replaces the string 'Cats in pajamas' with 'Cats in red pajamas'. Of course, we used a literal here for clarity. In actuality, 'Cats in pajamas' would be the value of a variable, and that value would be changed. Were this a function rather than a procedure, the value it would return would not be Cats in red pajamas' but 1, indicating that a single substitution had been made.

- The flag 'g' indicates that all matching portions of the target are to be replaced by the regular expression. Otherwise, only the first match is replaced.
- The way to distinguish the versions is by whether the target is a string (version 1 or 2) or a multi_line (version 3 or 4). Within these pairs, odd numbers are functions, and even procedures.

Here are the formal descriptions of the OWA_PATTERN procedures and functions.

owa_pattern.match (version 1)

Syntax

owa_pattern.match(line, pat, flags)

Purpose

Enables programmers to search a string for a pattern using regular expressions.

Parameters

line	in	varchar2
pat	in	varchar2
flags	in	varchar2 DEFAULT NULL

Generates

Boolean indicating whether match was found.

owa_pattern.match (version 2)

Syntax

owa_pattern.match(line, pat, flags)

Purpose

Enables programmers to search a string for a pattern using regular expression

Parameters

line	in	varchar2
pat	in out	pattern
flags	in	varchar2 DEFAULT NULL

Generates

Boolean indicating whether match was found.

owa_pattern.match (version 3)

Syntax

owa_pattern.match(line, pat, backrefs, flags)

Purpose

Enables programmers to search a string for a pattern using regular expression

Parameters

line	in	varchar2
pat	in	varchar2
backrefs	out	owa_text.vc_arr
flags	in	varchar2 DEFAULT NULL

Generates

Boolean indicating whether match was found.

owa_pattern.match (version 4)

Syntax

owa_pattern.match(line, pat, backrefs, flags)

Purpose

Enables programmers to search a string for a pattern using regular expression

Parameters

line	in	varchar2
pat	in out	pattern
backrefs	out	owa_text.vc_arr
flags	in	varchar2 DEFAULT NULL

Generates

Boolean indicating whether match was found.

owa_pattern.match (version 5)

Syntax

owa_pattern.match(mline, pat, rlist, flags)

Purpose

Enables programmers to search a string for a pattern using regular expression

Parameters

mline	inowa_text.multi_line	
pat	in	varchar2

rlist	out	owa_text.row_list
flags	in	varchar2 DEFAULT NULL

Generates

Boolean indicating whether match was found.

owa_pattern.match (version 6)

Syntax

owa_pattern.match(mline, pat, rlist, flags)

Purpose

Enables programmers to search a string for a pattern using regular expression

Parameters

mline	in	owa_text.multi_line
pat	in out	pattern
rlist	out	owa_text.row_list
flags	in	varchar2 DEFAULT NULL

Generates

Boolean indicating whether match was found.

owa_pattern.amatch (version 1)

Syntax

owa_pattern.match(line, from_loc, pat, flags)

Purpose

Enables programmers to search a string for a pattern using regular expressions.

Parameters

line	in	varchar2
from_loc	ini	nteger
pat	in	varchar2
flags	in	varchar2 DEFAULT NULL

Generates

Location (in number of characters from the beginning) of the end of the match. 0 if none such.

owa_pattern.amatch (version 2)

Syntax

owa_pattern.match(line, from_loc, pat, flags)

Purpose

Enables programmers to search a string for a pattern using regular expression

Parameters

line	in	varchar2
from_loc	in	integer
pat	in out	pattern
flags	in	varchar2 DEFAULT NULL

Generates

Location (in number of characters from the beginning) of the end of the match. 0 if none such.

owa_pattern.amatch (version 3)

Syntax

owa_pattern.match(line, from_loc, pat, backrefs, flags)

Purpose

Enables programmers to search a string for a pattern using regular expression

Parameters

line	in	varchar2
from_loc	in	integer
pat	in	varchar2
backrefs	out	owa_text.vc_arr
flags	in	varchar2 DEFAULT NULL

Generates

Location (in number of characters from the beginning) of the end of the match. 0 if none such.

owa_pattern.amatch (version 4)

Syntax

owa_pattern.match(line, from_loc, pat, backrefs, flags)

Purpose

Enables programmers to search a string for a pattern using regular expression

Parameters

line	in	varchar2
from_loc	in	integer
pat	in out	pattern
backrefs	out	owa_text.vc_arr
flags	in	varchar2 DEFAULT NULL

Generates

Location (in number of characters from the beginning) of the end of the match. 0 if none such.

owa_pattern.change (version 1)

Syntax

owa_pattern.match(line, from_str, to_str, flags)

Purpose

This version is a function. It enables programmers to search a string for a pattern and replace it.

Parameters

line	in out	varchar2
from_str	in	varchar2
to_str	in	varchar2
flags	in	varchar2 DEFAULT NULL

Generates

Revises line parameter. Function outputs number of substitutions made.

owa_pattern.change (version 2)

Syntax

owa_pattern.match(line, from_str, to_str, flags)

Purpose

This version is a procedure. It enables programmers to search a string for a pattern and replace it.



Parameters

line	in out	varchar2
from_str	in	varchar2
to_srt	in	varchar2
flags	in	varchar2 DEFAULT NULL

Generates

Revises line parameter.

owa_pattern.change (version 3)

Syntax

owa_pattern.match(mline, from_str, to_str, backrefs, flags)

Purpose

This version is a function. It enables programmers to search a multi_line for a pattern and replace it.

Parameters

mline	in out	multi_line	
from_str	in	varchar2	
to_srt	in	varchar2	
flags	in	varchar2	DEFAULT NULL

Generates

Revises mline parameter. Function outputs number of substitutions made.

owa_pattern.change (version 4)

Syntax

owa_pattern.match(mline, from_str, to_str, flags)

Purpose

This version is a procedure. It enables programmers to search a multi_line for a pattern and replace it.

Parameters

mline	in out	multi_line
from_str	in	varchar2

to_srt	in	varchar2
flags	in	varchar2 DEFAULT NULL

Generates

Revises mline parameter.

owa_pattern.getpat

Syntax

owa_pattern.getpat(arg, pat)

Purpose

This converts a VARCHAR2 string into the special datatype pattern. This datatypes is explained under OWA_PATTERN Datatypes.

Parameters

arg	in	varchar2
pat	in out	pattern

Generates

pattern.

OWA_TEXT

The OWA_TEXT package is chiefly used by OWA_PATTERN, but the functions are externalized so that you can use them directly if desired.

Datatypes

vc_array

A PL/SQL table of 32K VARCHAR2 strings, indexed by BINARY INTEGER. This is a component of *owa_text.multi_line.*

multi_line

A record with the following three fields:

int_array

A PL/SQL table of INTEGER indexed by BINARY INTEGER. This is a component of *owa_text.row_list.*

ORCL000930

row_list

A record with the following fields:

Procedures and Functions

owa_text.stream2multi

Syntax

owa_text.stream2multistream, mline)

Purpose

Converts a long string to a multi_line.

Parameters

srtream	in	varchar2
mline	out	multi_line

Generates

multi_line

owa_text.add2multi

Syntax

owa_text.add2multi(stream, mline, continue)

Purpose

Adds more content to a multi_line. The continue parameter specifies whether to begin appending within the previous final chunk (assuming it is less than 32K) or to start a new chunk.

Parameters

stream	in	varchar2
mline	out	multi_line
continue	in	boolean DEFAULT TRUE

Generates

multi_line

owa_text.new_row_list

Syntax

owa_text.new_row_list(rlist)

Purpose

Can be a procedure or a function. If a function it takes no parameters and outputs a new, initially empty, row_list. If a procedure, it places into the rlist parameter a new, initially empty, row_list.

Parameters

rlist	in	row_list

Generates

Actual output generated by htp.print.

owa_text.print_multi

Syntax

owa_text.print_multi(mline)

Purpose

Uses htp.print to print the multi_line.

Parameters

mline	in	multi_line

Generates

Actual output generated by htp.print.

owa_text.print_row_list

Syntax

owa_text.print_row_list(rlist)

Purpose

Uses htp.print to print the row_list.

Parameters

rlist	in	row_list

Generates

Actual output generated by htp.print.

OWA_IMAGE

OWA_IMAGE is a package providing simple functions that are for handling image maps.

Datatypes

point

This provides the x and y coordinates of a user's click on an image map..

Package Variables

null_point

This is a variable used to default point parameters. X and Y both are NULL..

Procedures and Functions

owa_image.get_x

Syntax

owa_image.get_x(p)

Purpose

This is a function that produces the X coordinate of the image map.

Parameters

p	in	point

Generates

X coordinate as integer.

owa_image.get_y

Syntax

owa_image.get_Y(p)

Purpose

This is a function that produces the Y coordinate of the image map.

Parameters

p in point

Generates

Y coordinate as integer.

OWA_COOKIE

OWA_COOKIE is a package that provides wrappers so that you can send cookies to and get them from the client's browser. Cookies are strings that are opaque to the client, but that maintain state throughout the client's session, or longer if an expiration date is included. Your system date will be calculated with reference to the information specified in the OWA_INIT package.

Datatypes

cookie

Since the HTTP standard is that cookie names can be overloaded, that is, multiple values can be associated with the same cookie name, this is a PL/SQL RECORD holding all values associated with a given cookie name. The fields are as follows:

```
name      varchar2(4K)
vals       vc_arr
numvals    integer
```

Note: vc_arr is defined in the OWA_TEXT package.

Procedures and Functions

owa_cookie.send

Syntax

owa_cookie.send(name, value, expires, path, domain, secure)

Purpose

This is a procedure that transmits a cookie to the client. This procedure must occur in the context of an OWA procedure's HTTP header output.

Parameters

name	in	varchar2
value	in	varchar2
expires	in	date DEFAULT NULL
path	in	varchar2 DEFAULT NULL
domain	in	varchar2 DEFAULT NULL
secure	in	varchar2 DEFAULT NULL

Generates

```
Set-Cookie: <name>=<value> expires=<expires> path=<path> domain=<domain>
secure
```

owa_cookie.get

Syntax

owa_cookie.get(name)

Purpose

This is a function that converts the string to a cookie

Parameters

name	in	varchar2

Generates

A cookie

owa_cookie.get_all

Syntax

owa_cookie.get_all(names, vals, num_vals)

Purpose

This is a procedure that returns all cookie name/value pairs from the client's browser that are associated with your domain.

Parameters

names	out	vc_arr
vals	out	vc_arr
num_valsout	integer	

Generates

Arrays of the names and of the values in the order received, and the count of the combinations.

owa_cookie.remove

Syntax

owa_cookie.remove(name, value, path)

Purpose

This is a procedure that forces a cookie to expire immediately. This output of this procedure must be embedded in an HTML header.

Parameters

name	in	varchar2
value	in	varchar2
path	in	varchar2 DEFAULT NULL

Generates

```
Set-Cookie: <name>=<value> expires=01-JAN-1990 path=<path>
```

OWA_INIT

This package provides constants that you override to set the time zone used by cookies. Cookies use expiration dates defined in Greenwich Mean Time (GMT)). If you are not on GMT, you can specify your time zone using one of these two constants. If your time zone is recognized by Oracle, which you can determine by consulting Chapter 3 of the *Oracle7 Server SQL Reference* under "SQL Functions", you can directly specify it using *dbms_server_timezone*; otherwise, you can use *dbms_server_gmtdiff* to give the offset of your time zone from GMT.

Change the values shown below and reload the package.

Constants

dbms_server_timezone *constant varchar2(3) := 'PST'*

Text string abbreviation for your timezone.

dbms_server_gmtdiff *constant number* *:= NULL*

Number of hours your timezone diverges from Greenwich Mean Time. A positive number if your timezone is ahead of GMT; otherwise negative.

CHAPTER

7 Oracle WebServer Messages

This appendix contains messages you may receive if you encounter errors while using Oracle WebServer. If you get an error message, look in this appendix to find the probable cause and suggestions of what to do.

00001 - 00600 Generic Oracle WebServer Configuration Messages

OWS—00000, "*utility* : normal, successful completion"

Cause: Normal exit

Action: None

OWS—00001, "*utility* : unable to open file *filename*"

Cause: The file may not exist or may not be readable to the signalling program.

Action: Verify the file's existence in the file-system. Check file and directory permissions and ownership, as well as the

ownership of the signalling process.

OWS—00002, "*utility* : error reading file *filename*"

Cause: Usually an operating system error. If the problem had been file permissions or non-existence, error 1 would have been signalled.

Action: Check your operating system log for errors.

OWS—00003, "*utility* : error writing file *filename*"

Cause: Usually an operating system error. If the problem had been file permissions or non-existence, error 1 would have been signalled.

Action: Check your operating system log for errors.

OWS—00010, "*utility* : file format error when reading *filename*"

Cause: The signalling program expected a specific file format, which it did not find.

Action: Analyze the file for incorrect formatting.

OWS—00011, "Installation not completed. Root.sh must be run"

Cause: The signalling program checked if root.sh completed successfully.

Action: Run root.sh as indicated by your installation guide.

OWS—00020, "*utility* : out of memory when requesting *number* bytes"

Cause: The operating system was not able to supply the requested amount of memory.

Action: Reduce the number of running programs on the machine and retry the operation.

OWS—00600, "*utility* : internal error, arguments [*arg1*][*arg2*][*arg3*][*arg4*][*arg5*]"

Cause: This is the generic internal error number for NDW program exceptions. This indicates that a process has encountered an exceptional condition.

Action: Report as a bug - the first argument is the internal error number

5000 - 5499 : Oracle Web Agent errors.

OWS—05100, "Agent : unable to connect due to Oracle error *number*"

Cause: The Oracle Web Agent was unable to connect to the specified Database due to an Oracle error.

Action: See the corresponding Oracle error. Common errors are: invalid username/password - check the Web Agent configuration file; Oracle not available - start up the database.

OWS—05101, "Agent : execution failed due to Oracle error *number*"

Action: See the Oracle error on the error stack.

OWS—05102, "Agent : initialization failed due to error *number*"

Action: See the Oracle error on the error stack.

OWS—05103, "Agent : error *number* during initialization"

Action: See the Oracle error on the error stack.

OWS—05104, "Agent : unable to attach to host database: [*ORACLE_HOME*], [*ORACLE_SID*], [*SQL*Net V2 Service*]"

Action: See the Oracle error on the error stack.

OWS—05110, "Agent : no stored procedure specified to call"

Cause: OWA requires the PATH_INFO environment variable to be set. PATH_INFO is set based on information passed through the URL after the OWA executable name and before the "?" or end of the URL

Action: specify a stored procedure name after the OWA executable name in the URL.

OWS—05111, "Agent : no stored procedure matches this call with the arguments passed"

Cause: The procedure and parameters passed to OWA don't match any existing stored procedures.

Action: verify that there is a stored procedure (not function) with no RECORD parameters which matches the parameters passed. If no value was passed for a parameter, be sure that the stored procedure was declared to have a default value for that parameter. Check for typographical errors.

OWS—05112, "Agent : too many procedures matches this call"

Cause: The procedure and parameters passed to OWA match too many existing stored procedures.

Action: This should only happen if there are overloaded stored procedures, and that in two different versions of the procedure, a given parameter is declared in one as a scalar and the other as a PL/SQL table.

OWS—05150, "Agent : http server Error - environment variable *variable* is NULL or non-existent"

Cause: The http server or other invoking application did not set the listed environment variable.

Action: Contact your http server or other invoking application vendor.

OWS—05151, "Agent : Port *number* not in list of valid ports : *number-list*"

Cause: The network port that the invoking http server is listening on is not a valid port for this OWA Service.

Action: Add this port to the list of valid ports for this Service, or access this service through a http server listening on a valid port.

OWS—05200, "Agent : service *service-name* not found in configuration file"

Cause: The piece of the URL just before the Web Agent executable name did not correspond to a Web Agent Service in the OWA configuration file.

Action: Correct the URL, or create the desired service.

OWS—05201, "Agent : service parameter *parameter* not found for service *service-name*"

Cause: A parameter required by the Oracle Web Agent was not found in the configuration file along with the specified service.

Action: Modify the Web Agent service configuration using the Web Agent administration form.

5500 - 5599 : Oracle Web Agent Administration errors and messages

OWS—05504, "Service *service-name* successfully written. Verify the PL/SQL output in *filename*"

Cause: The Web Agent Administration program successfully wrote the specified service and performed all actions requested. The program, however, cannot detect errors which may have occurred while installing the PL/SQL.

Action: Check the output in the file specified and verify that there are not unexpected error messages.

OWS—05521, "Service *service-name* submission failed because the passwords submitted do not match"

Cause: The password and the "confirm password" field do not match.

Action: Retype both passwords and re-submit the form.

OWS—05522, "Service *service-name* submission failed because password identification was selected and the submitted password was NULL"

Cause: The form submitted indicated that the Web Agent service should use a password to connect to the database, however no password was provided.

Action: Select operating system authentication for this service, or provide a password for the give OWA Database User.

OWS—05523, "Service *service-name* submission failed because value for mandatory parameter *parameter* is NULL"

Action: Fill in the specified field and re-submit the form.

OWS—05524, "Service *service-name* submission failed because both *parameter1* and *parameter2* are NULL"

Cause: One of the two parameters listed must not be null.

Action: Fill in one of the listed fields and re-submit the form.

OWS—05525, "Service *service-name* submission failed because a service with the same name already exists"

Action: Choose another name for the service you are attempting to create, or delete the existing service.

OWS—05526, "Service *service-name* submission failed due to error *number*"

Action: See the associated error and take the appropriate action.

OWS—05527, "Service *service-name* submission failed because the *PL/SQLAgent-or-DBA* Username and Password are invalid"

Cause: Either the OWA Database User name and password, or the DBA username and password, is not a valid username/password combination for the specified database.

Action: Correct the username and password and re-submit the form.

OWS—05528, "Service *service-name* submission failed because the DBA User submitted does not have *privilege* privileges"

Action: Choose a different Oracle username and password for the DBA user submitted, or have the specified privilege granted to this user by another DBA.

OWS—05529, "Service *service-name* to modify does not exist"

Cause: The user chose to modify an existing Web Agent service, however that service does not exist. This error should not occur unless the service to modify was deleted from a different form after this service modification form was generated.

Action: Create a new service.

OWS—05530, "Service *service-name* submission failed because the NLS Language could not be determined"

Cause: The Web Agent Administration package was unable to lookup the NLS Language for the specified database.

Action: Specify the value explicitly or submit a valid dba username and password.

ORCL000943

OWS—05531, "Service *service-name* submission failed because neither Server Manager, nor SQL*DBA could be run"

Cause: The Web Agent Administration package was unable to install the Oracle WebServer PS/SQL Web Toolkit which must be done from Server Manager (line mode) or SQL*DBA.

Action: Install either Server Manager or SQL*DBA in the ORACLE_HOME where the Web Agent administration is installed.

5600 - 5699 : Oracle Web Database Administration errors

OWS—05610, "Startup of database *database* failed due to error *number*"

Action: See the associated error and take the appropriate action.

OWS—05611, "Shutdown of database *database* failed due to error *number*"

Action: See the associated error and take the appropriate action.

OWS—05620, "DB Admin : no database selected"

Cause: A button was pressed to startup or shutdown the database, but no database had been selected.

Action: Return to the administration screen and select a database.

OWS—05621, "DB Admin : *database* is already running."

Cause: A request was made to start up a database, but that database is already running.

Action: If you would like to restart the database, shut it down first or choose "Startup Force".

OWS—05622, "DB Admin : *database* is already shut down."

Cause: A request was made to shut down a database, but that database is already shut down.

OWS—05623, "DB Admin : parameter file *filename* is larger than maximum size *number* bytes"

Cause: The initialization file listed was larger than the maximum allowable size.

Action: Reduce the size of the initialization file.

5700 - 5799 : Oracle Web Listener Configuration errors and messages

OWS—05710, "Value submitted for *parameter* must not be NULL"

Action: Type a value for the specified parameter and re-submit the form.

OWS—05711, "Port *number* is already in use by the Oracle Web Administration Server"

Cause: The port specified for service creation or modification is already in use by the Oracle WebServer Administration Server.

Action: Choose another port to run one of the Web Listeners on.

OWS—05712, "The Oracle Web Listener *listener-name* is already configured to run on port *number*"

Cause: The port specified for service creation or modification is already in use by another Oracle Web Listener.

Action: Choose another port to run one of the Web Listeners on.

OWS—05713, "Value for *parameter* must be between *minimum* and *maximum*"

Action: Specify a value between the two limits and resubmit the form.

OWS—05714, "For ports less than *number*, the effective userid must be root"

Cause: Ports below that specified require superuser privileges to run a Web Listener on it.

Action: Configure the Web Listener to run on a unrestricted port, or set the effective userid for this Web Listener the superuser.

OWS—05715, "An Oracle Web Listener named *listener-name* already exists"

Cause: Listener names must be unique

Action: Name the submitted listener something else.

OWS—05716, "Port Number *number* is duplicated in the Addresses and Ports list"

Cause: Port numbers must be unique

Action: Remove duplication and resubmit the form.

OWS—05717, "Fill out form has unknown input field format"

Cause: Oracle admbin/*.html files have probably been tampered with.

Action: Restore original files.

OWS—05721, "The Web Listener *listener-name* failed to start:"

Action: See the associated error message.

OWS—05722, "The Web Listener *listener-name* failed to reload:"

Action: See the associated error message.

OWS—05723, "The Web Listener *listener-name* was not stopped:"

Action: See the associated error message.

OWS-05725, "WRB Configuration file not copied from *pathname1* to *pathname2*"

Cause: Copying of file failed due to system error

Action: Make sure wladmin utility has permission to read/write.

OWS-05726, "Listener Name *listener-name* must be an alpha-numeric string"

Action: Only alphanumneric strings are accepted for listener name

OWS-05727, "Listener Name *listener-name* too long."

Cause: The specified Listener name is greater than six characters.

Action: Use a shorter name.

5800 - 5899 : Oracle Web Request Broker Administration messages

OWS-05801, "The input from the form in the HTML file *filename* has incorrect format"

Cause: Form has unexpected input field/value.

Action: Make sure only the OWS HTML templates are used to make submissions to the Web Server Manager utilities

OWS-05802, "The input parameter *parameter* must not be empty"

Cause: Required field not filled in submitted form.

Action: Fill in all required fields.

OWS-05804, "The WebServer Manager found no record of the listener *listener-name*"

Cause: Non-existent listener attempted to be accessed. owl.cfg might have been corrupted

Action: Make sure a non-existent listener is not referred to.

OWS-05805, "Could not copy WRB configuration from listener *listener-name1* to listener *listener-name2*"

Cause: WRB configuration file of "from" listener may not exist. System file copying error.

Action: Make sure wrbadmin utility has permissions to read and write, and the source file exists.

OWS-05810, "Configuration of cartridge *cartridge-name* for listener *listener-name* not found"

Cause: Attempt to modify or delete a cartridge configuration that doesn't exist

Action: Reload the WRB Configuration form to get current list of cartridges for this listener.

5900 - 5999 : Oracle WebServer Registration errors and messages

OWS—05901, "Please enter a value for field *field*."

Cause: The specified field is a required field for registration.

Action: Type a value into the specified field and re-submit the form.

7500 - 7599 : Oracle Web Server Proxy errors and messages

OWS—07500, "Proxy switch could not be accessed."

Cause: Error while reading or writing proxy switch value

Action: Check configuration file location/format.

OWS—07501, "Proxy Administration Form has invalid input values"

Cause: Form is being submitted with invalid field values

Action: Check and correct field values

Oracle Java Web Toolkit Messages

Message: [*Method*]: Invalid [*Argument*] Value(s)

Exception: `oracle.lang.HtmlRuntimeException`

Cause: The client programmer has supplied incorrect argument value(s) for the named [*Method*].

Action: Change the argument values.

Message: [*Method*]: Circular Reference

Exception: `oracle.lang.HtmlRuntimeException`

Cause: The client programmer has attempted to add a Container/ CompoundItem to itself.

Action: Remove that statement from the client application program.

Message: [*Method*]: File ([*filename*]) Not Found

Exception: `oracle.lang.HtmlRuntimeException`

Cause: This package/method cannot find the named file in the system's path.

Action: Provide the correct file/path name and/or change the system path.

Message: [*Method*]: IO Exception caught

Exception: `oracle.lang.HtmlRuntimeException`

Cause: An IO Exception has occurred.

Action: Check if the disk storage device is full or if there is a sharing violation.

APPENDIX

A Glossary

Administration DCD

The Database Connection Descriptor (DCD) used by the administration server to manage PL/SQL Agent database access.

Administration Server

A collection of special instances of WebServer components that a WebServer administrator uses to configure and maintain the WebServer.

Applet

A Java term for small programs that can be dynamically imported into Web pages or applications as needed. Generally, applets are imported from the Internet or another computer network.

Application Developer

A person who writes programs that the WebServer and/or the Oracle7 Server executes.

Authentication

a. A security scheme that requires a client to enter a user name and password to access certain files provided by the WebServer.

b. The practice of "signing" an electronic document in a legally binding way using digital signatures.

Base Directory

The directory name to which URL-encoded pathnames addressed to this

port are to be appended. For example, if the base directory is `/public_html`, the URL `http://www.blob.com/file` is converted to `http://www.blob.com/public_html/file`.

Basic Authentication

An authentication scheme that does not encrypt passwords when sending them over the Internet. Basic authentication is much less secure than digest authentication. See also *Digest Authentication* and *Authentication*.

Bytecode

Java code is interpreted in two steps. First, it is converted from source code (the Java code as written) to bytecode, which is a form executable by any platform on which Java runs. When the bytecode is executed, it is converted from bytecode to the native code for the platform in question.

CA

See *Certifying Authority (CA)*.

Certificate

A file provided by a certifying authority (CA) and installed on a WebServer machine that the WebServer uses to authenticate itself to clients requesting secure connections.

Certifying Authority (CA)

A trusted third-party company that provides certificates to legitimate organizations that request them.

CGI

See *Common Gateway Interface (CGI)*.

Character Set

A set of characters used to write a human language or group of languages, as defined by RFC 1521.

Client

A process, such as a Web browser, that interfaces to one or more users, sends requests to a server, and presents the results of those requests to the users. See also *Server*.

Common Gateway Interface (CGI)

The industry-standard technique for running applications on a web server. Oracle WebServer supports this standard, but also offers the Web Request Broker as a superior alternative.

Common Logfile Format

An industry standard format for transaction log files. The Web Listener uses this format to log transactions.

Configuration Directory

A directory in which a Web Listener process stores its configuration file.

Cookie

Information inserted by the server into the client's browser to track what the client has been doing. This can either expire when the user exits the browser or expire at the date the creator of the cookie specifies.

Database

A structured collection of information and the program that manages such. The Oracle7 Server is a relational database, which is the prevalent type.

Database Connection Descriptor (DCD)

A file that specifies information such as how the PL/SQL Agent is to connect to the Oracle7 Server to fulfil an HTTP request. A URL that requires the use of the PL/SQL Agent includes the name of the DCD it is to use. The information in the DCD includes the username (which also specifies the schema and the privileges), password, connect-string, error log file, standard error message, and the language to be used.

Data Integrity

A mechanism that uses digital signatures to ensure that transmitted data is not tampered with.

DCD

See *Database Connection Descriptor (DCD)*.

Default Character Set

The character set the Web Listener uses in interpreting a file that uses an unrecognized character set.

Default DCD

The Database Connection Descriptor(DCD) that the PL/SQL Agent uses to access a database in response to a request that does not specify an available DCD.

Default MIME Type

The Multipurpose Internet Mail Extensions (MIME) type that the Web Listener uses in interpreting requested files of an unsupported MIME type. See also *MIME Type*.

Digest Authentication

An authentication scheme that encrypts passwords before sending them over the Internet, unlike basic authentication. See also *Authentication*.

Digital Signature

A "signature" attached to an electronic document that reliably identifies the author or sender, and guarantees that the document has not been tampered with.

Directory Indexing

The practice of returning a directory listing when a request URL resolves to a directory that does not contain the default initial file. Directory listings can sometimes help clients learn the correct spelling of filenames they want to request.

Directory Mapping

The practice of defining a virtual file system.

Dispatcher

See *WRB Dispatcher*.

DNS Resolution

The practice of determining a computer's DNS (Domain Name Service) host name from its IP address.

Document Root

The file-system directory that serves as the root of the Web Listener's virtual file system.

Domain

See *Domain Name Service (DNS)*.

Domain-based Restriction

A restriction scheme that allows only machines within specified DNS domains to access certain files. See also *Restriction*.

Domain Name Service (DNS)

The mechanism that divides the Internet into separate, hierarchical groups called domains, identified by unique alphanumeric names, such as `us.oracle.com`. DNS identifies each computer within a domain by a unique host name. For example, a computer named `hal` in the `us.oracle.com` domain would be uniquely identified on the Internet as `hal.us.oracle.com`

Encoding

An algorithm used to alter a file's format, such as compression. You can use the WebServer manager to define the encodings that each Web Listener process recognizes.

Encryption

The practice of scrambling (encrypting) data in such a way that only an intended recipient can unscramble (decrypt) and read the data. See also *Public-Key Encryption* and *Secret-Key Encryption*.

Error File

A file to which a Web Listener process logs errors. There is one error file for each Web Listener process.

Exception

A runtime occurrence in PL/SQL or Java that requires special handling and may indicate an error.

File Caching

The practice of leaving files open (resident in memory) so the WebServer can provide them to clients quickly. You can use the WebServer Manager to specify files to be cached.

Filename Extension

A short alphanumeric suffix attached to a filename, following a dot "." that represents the file's format. The WebServer uses filename extensions to identify several kinds of file formats, including MIME types and encodings. A file may have several extensions.

File Protection

The practice of assigning an authentication or restriction scheme to control access to a specific file or group of files.

Firewall Machine

A computer that regulates access to computers on a local area network from outside, and regulates access to outside computers from within the local area network.

Foreign Key

See *Key*.

genreq

A utility you can use to generate a request for a certificate. You can then submit the generated request to a certifying authority (CA).

Host Name

An alphanumeric character string that uniquely identifies a computer within a DNS domain.

HTTP

See *HyperText Transfer Protocol (HTTP)*.

HTTPS

Secure HTTP—a version of HTTP with provisions for secure data transmission. See *HyperText Transfer Protocol (HTTP)*.

HyperText Markup Language (HTML)

A format for encoding hypertext documents that may contain text, graphics, and references to programs, and references to other hypertext documents.

HyperText Transfer Protocol (HTTP)

The protocol that clients use to issue requests for documents to the WebServer.

Image Map

Graphic in a Web page that specifies several URL's, each associated with a specified region of the single image.

Info File

A file to which a Web Listener process logs its transactions on a particular port. There is one info file for each port on which the Web Listener process accepts connections. The info file is in Common Logfile Format.

Initial File

The name of the HTML file that the WebServer returns by default when a request URL specifies only a directory name.

IP Address

A four-part number with no more than three digits in each part that uniquely identifies a computer on the Internet; the number format is defined by the Internet Protocol (IP).

IP-based Restriction

A restriction scheme that allows only machines within specified groups of IP addresses to access certain files. See also *Restriction*.

Java

Language developed by Sun Microsystems and used by Oracle WebServer. This language is fully object-oriented, extremely portable, and

optimized for creating distributed applications on the Internet or other computer networks. Oracle WebServer can execute Java directly and can send Java programs called applets to the client's browser for execution there.

Java Interpreter

In general, a program that interprets and executes Java bytecode independently of a Web browser. In Oracle WebServer, this refers specifically to the Java Interpreter that WebServer provides for the purpose of executing Java on the server.

Java Web Developer's Toolkit

A group of Java classes provided with the Oracle WebServer SDK to make it easier for you to interface to the Oracle7 Server and generate dynamic HTML using Java.

Key

a. A large number used in encrypting data. See also *Private Key* and *Public Key*.

b. A unique identifier used in a relational database, called a "primary key."

c. A reference to a primary key, called a "foreign key."

Key Pair

A pair of mathematically related keys (a public key and a private key) associated with a user, used in public-key encryption.

Language Identifier

A two-character alphanumeric string that identifies a human language, as defined by RFC 1766.

Listener

Portion of the WebServer that receives HTTP requests. You can use the Oracle WebServer Manager to create multiple Web Listener processes and assign each to accept connections on a different set of ports.

Listener Configuration

A collective name for the parameters that control the behavior of a Web Listener process. You use the WebServer Manager to maintain listener configurations.

Listener Name

An alphanumeric string no more than six characters long that uniquely identifies a Web Listener process.

Listener PID File

A one-line text file that contains an ASCII representation of the process ID (PID) of a Web Listener process.

LiveHTML

Oracle's extension of the industry-standard Server Side Includes (SSI) functionality. LiveHTML files supplement HTML with instructions that WebServer executes before transmitting the page. These instruction specify material that is to be included in the generated page. Said material can include other Web pages, environment variables, and the output of programs executed on the Server. The programs may, but need not, conform to the CGI standard.

Local Database

For the WebServer, a database that runs on the same machine as the WebServer.

Log File Directory

A directory in which the PL/SQL Agent stores log files for a particular Database Connection Descriptor (DCD).

Memory Mapping

For the WebServer, the practice of mapping an open file directly into the address space of a Web Listener process. This speeds file access, and allows multiple clients to access the same file simultaneously without making a separate copy for each client.

MIME Type

A file format defined by the Multipurpose Internet Mail Extensions standard. Several RFCs define MIME.

NULL

a. A marker is the database for the absence of data.

b. The logical result of the comparison of a database NULL with any value.

c. A statement is PL/SQL that does nothing, but functions as a placeholder.

Oracle7 Server

The leading database product in the world. A program for sophisticated high-level management of information. See also *Database* and *Server*.

ORACLE_HOME

Environment Variable that indicates the root of the Oracle7 Server code

tree.

Oracle Web Agent (OWA)

A term from an earlier release of this product. OWA is now the PL/SQL Agent. For the sake of compatibility, the string "owa" is still used in URLs to specify that execution is to go to the PL/SQL Agent.

Overloading

Overloaded procedures and functions (in PL/SQL) or methods (in Java) have the same name but take different parameters (PL/SQL) or are contained in different classes (Java) and do similar but not identical things.

Package

A group of PL/SQL functions and procedures.

Parsable File

A file located on the WebServer that contains codes that the server interprets prior to transmission of the file as a Web page. This is part of the LiveHTML functionality. See also *LiveHTML*.

PL/SQL

Oracle's proprietary extension to the SQL language. PL/SQL adds procedural and other constructs to SQL that make it suitable for writing applications.

PL/SQL Agent

A server extension that interfaces to the Oracle7 Server using PL/SQL and dynamically derives and outputs HTML. The PL/SQL Agent can be invoked either as a WRB Service or as a CGI program.

PL/SQL Developer's Toolkit

A bundle of PL/SQL packages, provided with the Oracle WebServer SDK, that make it easier to generate HTML using PL/SQL. Applications written for either the PL/SQL Agent or the Java Interpreter can use these packages.

Port

A number that TCP uses to route transmitted data to and from a particular program.

Preferred Language

The language the Web Listener uses when handling a request for a file available in more than one language, if the request doesn't specify a language. Language identifiers are defined by RFC 1766.

Primary Key

See *Key*.

Private Key

A key known only to one user, used to decrypt data encrypted with the user's public key. See also *Public-Key Encryption*.

Proxy Server

An HTTP engine, such as the Web Listener, that clients inside a firewall can use to access web sites outside the firewall.

Public Key

A key known to all users, used to encrypt data in such a way that only a specific user can decrypt it. See also *Private Key* and *Public-Key Encryption*.

Public-Key Encryption

A form of encryption that uses a key pair (a public key and a private key) to encrypt and decrypt data.

Query String

Optional portion of a URL that specifies parameters to be passed to some server extension.

Realm

A group of users and groups assigned by an authentication scheme to regulate access to specific files or directories.

Remote Database

For the WebServer, a database running on a different machine from the WebServer, which the WebServer accesses over the network.

Restriction

A security scheme that restricts access to files provided by the WebServer to client machines within certain groups of IP addresses or DNS domains.

Routing

The process of directing data from one machine on the Internet to another by way of intermediate machines.

Secret-Key Encryption

A form of encryption that uses a single key both to encrypt and to decrypt a document. Secret-key encryption is much faster that public-key encryption, but is more vulnerable to attack.

Secure Sockets Layer (SSL)

An emerging standard for secure transmission of hypertext documents over the Internet using secure HTTP (HTTPS).

Server

A process that executes requests on behalf of another process (the client) whose main purpose is to interface to the user. There are two types of servers relevant to this product. The first is the Oracle7 Server, which is a *database* server dedicating to performing data management duties on behalf of clients using any number of possible interfaces. The other is the Oracle WebServer itself, which is a *web* server, dedicating to answering requests that come in through the HyperText Transfer Protocol (http). This product is a web server that utilizes a database server.

Server Extension

Generic term for external programs executed by the Oracle WebServer whether through the WRB or CGI. Server Extensions are either WRB Services or CGI programs.

Server-Parsable

See *Parsable File*.

Server Side Includes (SSI)

Industry-standard term for LiveHTML. See *LiveHTML*.

Session Key

A secret key used by SSL to encrypt data transmitted over a secure connection. The client generates the session key after the WebServer authenticates itself and communicates it to the WebServer using public-key encryption.

Socket

The combination of an IP address and a port number.

SQL (Structured Query Language)

The industry-standard language for interfacing to relational databases such as the Oracle7 Server. See also *PL/SQL*.

SSI

See *Server Side Includes (SSI)* and *LiveHTML*.

SSL

See *Secure Sockets Layer (SSL)*.

Table

a. In HTML, a table is a way of presenting information to the user.

b. In SQL, a table is the basic way that data is structured, regardless of how it is presented to the user.

c. In PL/SQL, special kinds of SQL tables are used to function, effectively, as dynamic arrays.

Transmission Control Protocol (TCP)

The underlying communication protocol that the WebServer and its clients use to communicate HTTP requests.

Uniform Resource Locator (URL)

The text-string format clients use to encode requests to the WebServer.

User Directory

The subdirectory of a user's home directory in which the Web Listener searches by default for files when the user's home directory appears in the request URL.

Varchar2

A standard datatype of the Oracle7 Server. A variable-length string.

Virtual Machine (VM)

The mechanism the Java language uses to achieve its high portability. Java bytecode is executable by any Java Virtual Machine running on any actual machine. The VM converts the bytecode to the native code for the machine at hand.

Virtual File System

A mapping that associates the pathnames used in request URLs to the file system maintained by the host machine's operating system.

Virtual Pathname

A synonym that the virtual file system maps to a file stored in the file system maintained by the host machine's operating system.

Web Listener

See *Listener*.

Web Request Broker (WRB)

The core of the Oracle WebServer architecture. The Web Request Broker passes http requests that require the running of Server programs to various processes (WRB Executable Engines or WRBX's) that continuously run and await such requests. The WRB also includes an open

API, so you can run your own Server programs under it. The WRB is a more efficient alternative to the industry-standard CGI, but it can process requests that use CGI environment variables.

WebServer Administrator

Person in charge of configuring and running the Oracle WebServer.

WebServer Manager

A collection of utilities and HTML forms you can use to configure and maintain the WebServer installed on your computer.

Wrapper

A Java class that encapsulates another kind of object, possibly external to Java itself, such as a PL/SQL package.

WRB

See *Web Request Broker (WRB)*.

WRB API

An open API used by the Web Request Broker (WRB).

WRB Cartridge

A program that is executed on the WebServer through the WRB API.

WRB Dispatcher

Portion of the WRB that distributes requests to running processes.

WRB Executable Engine (WRBX)

One of a pool of processes that the WRB maintains continuously, so that HTTP requests requiring the execution of programs are not slowed down by the performance cost of spawning a new process. WRBX's are associated with WRB Cartridges and are created and destroyed according to workload.

WRB Service

A particular server extension to be run through the WRB. A WRB Service is the combination of a WRB Cartridge with its interface, the WRB API. The Oracle WebServer comes with three WRB Services: the PL/SQL Agent, the Java Interpreter, and the LIveHTML Interpreter. Using the WRB API, you can supplement these with your own.



APPENDIX

B

Overview of the Oracle7 Server, SQL, and PL/SQL

This section provides an overview of the Oracle database, and a guide to its documentation set. This subject cannot be thoroughly covered here. What is provided is a conceptual overview, and enough information on PL/SQL to enable you to write simple programs. You are referred to the Oracle7 Server documentation for further information. The topics covered in this section are as follows:

- The *Oracle7 Server*. A conceptual overview of what a relational database is.
- *SQL*. An introduction to the standard language used to interface to relational databases.
- *PL/SQL*. An introduction to Oracle's extension of SQL that makes it a versatile high-level programming language.

Oracle7 Server

The Oracle7 Server is a Relational Database Management System (RDBMS). That is to say, the job of the Oracle7 Server is to manage data. Users or other processes store, alter, and obtain the data by issuing statements in *SQL* that the RDBMS executes, but they never directly access the data. Having all the information under the control of a single entity ensures, for example, that the information

maintains a coherent structure and that simultaneous changes by different users do not interfere with one another.

Database Tables

Saying that Oracle7 is a *Relational* Database Management System implies that all of the data it contains is structured as tables (tables are called "relations" in mathematical jargon).

Here is a simple table, such as you might find in an Oracle database:

CNUM	FNAME	*LNAME*	*ADDRESS*
4005	Julia	Peel	197 Myrtle Court, Brisbane, CA
4007	Terry	Subchak	2121 Oriole Way, Boston, MA
4008	Emilio	Lopez	31D San Bruno Ave. SF, CA
4011	Kerry	Lim	455 32nd St. #45, Brinton, KY

The Customers Table

Each row of this table describes one person, and each column has one type of information about that person. Note the column *cnum*. This is simply a number we generate to distinguish the customers from one another, as names are not necessarily unique. You refer to data in a relational database by its content, not by such things as where it is stored. Therefore, every table must have an identifying group of one or more columns whose values, taken as a set, are always different for every row of the table. This group, in this case the single column *cnum*, is called the *primary key* of the table. Locally generated numbers, as in this example, are a common and easy way to create primary keys.

Foreign Keys

Suppose you wanted to add our customers' phone numbers to your database. Since one person can have multiple phone numbers, these do not fit into your structure well. If you want to include the phone numbers in the Customers table, there are three possibilities, none of them good:

1. You could fit all the phone numbers for a given person into a single column in a single row, in which case it would be difficult to access the phone numbers independently.
2. You could create a column for each type of phone number, in which case you would have to redesign your table each time a new type arose. This also could create an unwieldy number of mostly empty columns.
3. You could enter a new row for each phone number, in which case each such row would consist of redundant information except for the new phone number. This approach would be error-prone and waste space.

The good solution is simply to create a second table, like this:

CNUM	PHONE	TYPE
4005	375-296-8226	home
4005	375-855-3778	beeper
4008	488-255-9011	home
4011	577-936-8554	home
4008	488-633-8591	work

Table 0 - 1 Customers_Phone

Notice that in this table *cnum* is not the primary key; it identifies the customer and therefore is the same for each phone number associated with a given customer. What, then, is the primary key? The combination of *cnum* and *phone*. If you list the same number for the same person twice, you really have made a duplicate entry and should eliminate one anyway.

The *cnum* column does have a special function, however, because it defines the relationship between Customers_Phone and Customers by associating each phone number with a customer. We say that it *references* the *cnum* column in Customers. A group of one or more columns, such as this, that references another group is known as a *foreign key*. The group of columns a foreign key references is called its *parent key* or its *referenced key*. Each foreign key value references a specific row in the table containing the parent key. Clearly, then, all sets of values in the foreign key have to be present once and only once in the parent key (although they may be present any number of times in the foreign key itself, as above) for the reference to be both meaningful and unambiguous. For that reason, the parent key must be either a primary key (the usual case) or another group of columns that is unique, which is known as a *unique key*.

Oracle can make sure that all primary and unique keys stay unique and that all foreign key references are valid; this is called maintaining *referential integrity*. For more information on foreign and parent keys, see Chapter 7 of the *Oracle7 Server Concepts Manual* and "CONSTRAINT clause" in Chapter 4 of the *Oracle7 SQL Reference*.

Users, Connections, Privileges, and Roles

The Oracle7 Server controls which users can do what to its data. To do this, it uses a log-on procedure that is separate from that of the operating system. Once logged on the operating system, you establish a *connection* to the Oracle7 Server with a username and password known to Oracle that may have no relationship

to the ones used for you by the operating system. For more information, see "CREATE USER" and "CONNECT" in the *Oracle7 Server SQL Reference*.

The name under which you connect to Oracle7—your Oracle username— is associated with a number of *privileges*, which are the rights to perform various actions. For more information on privileges, see "GRANT" and "REVOKE" in Chapter 4 of the *Oracle7 Server SQL Reference*.

The PL/SQL Agent as an Oracle User

When you access the Oracle7 Server through the WebServer, you use a PL/SQL Agent that already has an established connection to the Server. When a URL causes the Web Listener to invoke the PL/SQL Agent, it associates the PL/SQL Agent with a service based on the URL and/or the domain name of the issuer of the URL. This service determines the Oracle user that the PL/SQL Agent behaves as when executing this request, and thereby controls what that request can do.

SQL

SQL (Structured Query Language) is the language you use to issue instructions to the *Oracle7 Server*. It is, in fact, the standard language used by all major relational database vendors, and Oracle complies at Entry Level with SQL92, the most recent ISO (International Standards Organization) standard. There are several aspects of SQL that may differ from computer languages that you are familiar with, such as the following:

- SQL is non-procedural. In SQL, you tell the Server what to do but not how it is to be done. This frees you from dealing with a lot of detail.
- SQL statements are independent of one another. Although *PL/SQL* addresses this, SQL itself has no conditional or other control-flow statements.
- SQL employs set-at-a-time operation. It operates on arbitrarily large sets of data in a single step.
- SQL uses *Nulls and Three-Valued Logic*. In most languages, Boolean expressions are either TRUE or FALSE. In SQL, they are TRUE, FALSE, or NULL. This will be explained shortly.

Retrieving Data

Suppose you wanted to pull from the Customers table the information on customers named "Peel". This is called making a *query*. To do it, you could issue the following statement:

```
SELECT *
     FROM Customers
     WHERE LNAME = 'Peel';
```

This produces the following:

```
CNUM       FNAME              LNAME      ADDRESS

4005       Julia              Peel       197 Myrtle Court, Brisbane, CA
```

Oracle interprets the statement as follows: Any number of spaces and/or line breaks are equivalent to one space or line break. These are delimiters, and the extra spaces and line breaks are for readability: all are equivalent "white space". Likewise, case is not significant, except in literals like the string you are searching for ('Peel').

SELECT is a keyword telling the database that this is a query. All SQL statements begin with keywords. The asterisk means to retrieve all columns; alternatively, you could have listed the desired columns by name, separated by commas. The FROM Customers clause identifies the table from which you want to draw the data.

WHERE LNAME = 'Peel' is a predicate. When a SQL statement contains a predicate, Oracle tests the predicate against each row of the table and performs the action (in this case, SELECT) on all rows that make the predicate TRUE. This is an example of set-at-a-time operation. The predicate is optional, but in its absence the operation is performed on the entire table, so that, in this case, the entire table would have been retrieved. The semi-colon is the statement terminator.

Nulls and Three-Valued Logic

With predicates, you should be aware of three-valued logic. In SQL, the basic Boolean values of TRUE and FALSE are supplemented with another: NULL, also called UNKNOWN. This is because SQL acknowledges that data can be incomplete or inapplicable and that the truth value of a predicate may therefore not be knowable. Specifically, a column can contain a *null*, which means that there is no known applicable value. A comparison between two values using relational operators—for example, a = 5—normally is either TRUE or FALSE. Whenever nulls are compared to other values, however, including other nulls, the Boolean value is neither TRUE nor FALSE but itself NULL.

In most respects, NULL has the same effect as FALSE. The major exception is that, while NOT FALSE = TRUE, NOT NULL = NULL. In other words, if you know that an expression is FALSE, and you negate (take the opposite of) it, then you know that it is TRUE. If you do not know whether it is TRUE or FALSE, and you negate it, you still do not know. In certain cases, three-valued logic can create problems with your programming logic if you have not accounted for it. You can treat nulls specially in SQL with the IS NULL predicate, as explained in Chapter 3 of the *Oracle7 Server SQL Reference*.

Creating Tables

This is how you create tables in SQL. You can use the following SQL statement to create the Customers table:

```
CREATE TABLE Customers
(cnum           integer NOT NULL PRIMARY KEY,
FNAME           char(15) NOT NULL,
LNAME           char(15) NOT NULL,
ADDRESS          varchar2 );
```

After the keywords CREATE TABLE come the table's name and a parenthesized list of its columns with a definition of each. Integer, char, and varchar2 are datatypes: all of the data in a given column is always of the same type (char means a fixed and varchar2 a varying length string). For more information on SQL datatypes, see Chapter 2 of the *Oracle7 Server SQL Reference*.

NOT NULL and PRIMARY KEY are constraints on the columns they follow. They restrict the values you can enter in those columns. Specifically, NOT NULL forbids you from entering nulls in the column. PRIMARY KEY prevents you from entering duplicate values into the column and makes the column eligible to be the parent for some foreign key. For more information, see "CREATE TABLE" and "CONSTRAINT clause" in Chapter 4 of the *Oracle7 Server SQL Reference*.

Ownership and Naming Conventions

Note that when you create a table in SQL, you own it. This means you generally have control over who has access to it, and that it is part of a schema that bears your Oracle username. A *schema* is a named collection of database objects under the control of a single Oracle user. Schemas inherit the names of their owners. When other users refer to an object you have created, they have to precede its name by the schema name followed by a dot (no spaces). SQL utilizes a hierarchical naming convention with the levels of the hierarchy separated by dots. In fact, you sometimes have to precede column names by table names to avoid ambiguity, in which case you also use a dot. The following is an example in the form *schemaname.tablename.columnname:*

```
scott.Customers.LNAME
```

You can simplify references like this by using *synonyms*, which are aliases for tables or other database objects. Synonyms can be *private*, meaning that they are part of your schema and you control their usage, or *public*, meaning that all users can access them. For example, you can create a synonym "Cust" for scott.Customers as follows:

```
CREATE SYNONYM Cust FOR scott.Customers;
```

This would be a private synonym, which is the default. Now you could rewrite the example above like this:

```
Cust.LNAME
```

You still have to refer to the column directly. Synonyms can only be for tables, not table components like columns.

For more information on synonyms, see "CREATE SYNONYM" in Chapter 4 of the *Oracle7 Server SQL Reference*. For more information on SQL naming conventions, see Chapter 2 of the *Oracle7 Server SQL Reference*. For more on schemas, see "CREATE SCHEMA" in Chapter 4 of the *Oracle7 Server SQL Reference*.

Inserting and Manipulating the Data

Which SQL statements determine the actual data content? Chiefly, three —the INSERT statement, the UPDATE statement, and the DELETE statement. INSERT places rows in a table, UPDATE changes the values they contain, and DELETE removes them.

The INSERT Statement

For INSERT, you simply identify the table and its columns and list the values, as follows:

```
INSERT INTO Customers (cnum, FNAME, LNAME)
     VALUES (2004, 'Harry', 'Brighton');
```

This statement inserts a row with a value for every column but ADDRESS. Since you did not, in your CREATE TABLE statement, place a NOT NULL constraint on the ADDRESS column, and since you did not give that column a value here, Oracle sets this column to *null*. If you are inserting a value into every column of the table, and you have the values ordered as the columns are in the table, you can omit the column list. You optionally can put a SELECT statement in place of the VALUES clause of the INSERT statement to retrieve data from elsewhere in the database and duplicate it here. For more information on the INSERT and the

SELECT statements, see "INSERT" and "SELECT," respectively, in Chapter 4 of the *Oracle7 Server SQL Reference*.

The UPDATE Statement

UPDATE is similar to SELECT in that it takes a predicate and operates on all rows that make the predicate TRUE. For example:

```
UPDATE Customers
     SET ADDRESS = null
     WHERE LNAME = 'Subchak';
```

This sets to null all addresses for customers named 'Subchak'. The SET clause of an UPDATE command can refer to current column values. "Current" in this case means the values in the column before any changes were made by this statement. For more information on the UPDATE statement, see "UPDATE" in Chapter 4 of the *Oracle7 Server SQL Reference*.

The DELETE Statement

DELETE is quite similar to UPDATE. The following statement deletes all rows for customers named 'Subchak':

```
DELETE FROM Customers
     WHERE LNAME = 'Subchak';
```

You can only delete entire rows, not individual values. To do the latter, use UPDATE to set the values to null. Be careful with DELETE that you do not omit the predicate; this empties the table. For more information on DELETE, see "DELETE" in Chapter 4 of the *Oracle7 Server SQL Reference*.

Querying Multiple Tables Through Joins

Even though it only retrieves data, SELECT is the most complex statement in SQL. One reason for this is that you can use it to query any number of tables in one statement, correlating the data in various ways. One way to do this is with a *join*, which is a SELECT statement that correlates data from more than one table. A join finds every possible combination of rows, such that one row is taken from each table joined. This means that three tables of ten rows each can produce a thousand rows of output (10 * 10 * 10) when joined. Typically, you use the predicate to filter the output in terms of some relationship. The most common type of join, called a *natural join*, filters the output in terms of the foreign key/parent key relationship explained earlier in this appendix. For example, to see the people in the Customers table coupled with their various phone numbers from the Customers_Phone table, you could enter the following:

```
SELECT a.CNUM, LNAME, FNAME, PHONE, TYPE
     FROM Customers a, Customer_Phone b
     WHERE a.CNUM = b.CNUM;
```

In the above, a and b are *range variables*, also called *correlation variables*. They are simply alternate names for the tables whose names they follow in the FROM clause, so that a = Customers and b = Customers_Phone. You can see that here you need the range variables to distinguish Customers.CNUM from Customers_Phone.CNUM in the SELECT and WHERE clauses. Even when not needed, range variables are often convenient.

Here is the output of the natural join:

CNUM	LNAME	FNAME	PHONE	TYPE
4005	Peel	Julia	375-296-8226	home
4005	Peel	Julia	375-855-3778	beeper
4008	Lopez	Emilio	488-255-9011	home
4008	Lopez	Emilio	488-633-8591	work
4011	Lim	Kerry	577-936-8554	home

This output represents every combination of rows from the two tables where both rows have the same CNUM value.

Outer Joins

Notice in the preceding example that people from the Customers table who did not have phones (namely, CNUM 4007) were not selected. If a row has no match in the other table, the predicate is never true for that row. Sometimes, you do not want this effect, and you can override it by using an *outer join*. An outer join is a join that includes all of the rows from one of the tables joined, regardless of whether there were matches in the other table. Such a join inserts nulls in the output in whichever columns were taken from the table that failed to provide matches for the outer-joined table. Here is the same query done as an outer join:

```
SELECT a.CNUM, LNAME, FNAME, PHONE, TYPE
     FROM Customers a, Customer_Phone b
      WHERE a.CNUM = b.CNUM (+);
```

This is the output of the above:

CNUM	LNAME	FNAME	PHONE	TYPE
4005	Peel	Julia	375-296-8226	home
4005	Peel	Julia	375-855-3778	beeper
4007	Subchak	Terry	NULL	NULL
4008	Lopez	Emilio	488-255-9011	home
4008	Lopez	Emilio	488-633-8591	work
4011	Lim	Kerry	577-936-8554	home

Notice that the only difference in the query is the addition of (+) to the WHERE clause. This follows the table for which nulls are to be inserted. The output from the query, then, includes at least one row for each row of the table that *did not* have (+) appended in the predicate.

You can also use SELECT statements to produce values for processing within queries (these are called subqueries), and you can perform standard set operations (UNION, INTERSECTION) on SELECT statement output. For more information on the SELECT statement, subqueries, and joins, see "SELECT" in Chapter 4 of the *Oracle7 Server SQL Reference*.

Where to Look for More Information

Oracle7 SQL is a very complex subject, and we have been able only to scratch the surface of it here. To make it easier for you to find the specific information you need to perform the task at hand, we provide the following table, which identifies where in the Oracle7 Server documentation set you can find

information on specific SQL topics. Unless otherwise noted, find the headings in Chapter 4 of the *Oracle7 Server SQL Reference*.

To Find Out About	Look Under
aggregate data (totals, counts, averages, and so on)	SQL Functions in Chapter 3 of the *Oracle7 Server SQL Reference*.
changing user passwords	ALTER USER
connecting to the database	CONNECT
constraints	CONSTRAINT clause; CREATE TABLE; ENABLE clause
controlling user access to objects and user actions	GRANT; REVOKE; CREATE ROLE; SET ROLE; see also Chapters 17 and 18 in the *Oracle7 Server Concepts Manual*
creating databases	CREATE DATABASE
creating users	CREATE USER
functions that change simple values	SQL Functions in Chapter 3 of the *Oracle7 Server SQL Reference*.
linking databases at different locations	CREATE DATABASE LINK; see also "Distributed Databases" in the *Oracle7 Server Concepts Manual*.
making changes to the data permanent	COMMIT; SET TRANSACTION; SAVEPOINT
making SQL statements execute more quickly	CREATE INDEX; see also "Indexes" in the *Oracle7 Server Concepts Manual*.
monitoring database usage	AUDIT
reversing (undoing) changes to the data	ROLLBACK; SET TRANSACTION; SAVEPOINT

PL/SQL

PL/SQL is an application-development language that is a superset of *SQL*, supplementing it with standard programming-language features that include the following:

- block (modular) structure

- flow-control statements and loops
- variables, constants, and types
- structured data
- customized error handling

Another feature of PL/SQL is that it allows you to store compiled code directly in the database. This enables any number of applications or users to share the same functions and procedures. In fact, once a given block of code is loaded into memory, any number of users can use the same copy of it simultaneously (although behavior is as though each user had her own copy), which is useful for the Oracle WebServer. PL/SQL also enables you to define triggers, which are subprograms that the database executes automatically in response to specified events.

Unlike SQL, PL/SQL is not an industry standard, but is an exclusive product of Oracle Corporation.

The remainder of this section covers the following PL/SQL topics:

- *Basic Structure and Syntax*
- *The DECLARE Section*
- *The EXECUTABLE Section*
- *The EXCEPTION Section*
- *Storing Procedures and Functions in the Database*
- *Database Triggers*

Note: For the sake of efficiency, PL/SQL code is compiled prior to runtime. It cannot refer at compile time to objects that do not yet exist, and, for that reason, the one part of SQL that PL/SQL does not include is DDL (Data Definition Language)—the statements, such as CREATE TABLE, that create the database and the objects it contains. However, you can work around this by using the package DBMS_SQL, included with the server, to generate the DDL code itself dynamically at runtime. For more information, see "Using DDL and Dynamic SQL" in the *PL/SQL User's Guide and Reference*.

Basic Structure and Syntax

PL/SQL, like many programming languages, groups statements into units called *blocks*. These can either be named, in which case they are called *subprograms*, or unnamed, in which case they are *anonymous blocks*. Subprograms can be either functions or procedures. The difference between these, as in most languages, is that a function is used in an expression and returns a value to that expression, while a procedure is invoked as a standalone statement and passes values to the calling program only through parameters. Subprograms can be nested within one another and can be grouped in larger units called *packages*.

A block has three parts:

- *The DECLARE Section*. This is where you define local variables, constants, types, exceptions, and nested subprograms. PL/SQL has a forward declaration, but you can use it only for subprograms. Therefore, you must define all variables, constants, and types before referencing them. For more information on forward declarations, see "Declaring Subprograms" in the *PL/SQL User's Guide and Reference*.
- *The EXECUTABLE Section*. This is the actual code that the block executes. This is the only part of the block that must always be present.
- *The EXCEPTION Section*. This is a section for handling runtime errors and warnings.

These divisions are explained further in the sections that follow.

The DECLARE Section

The DECLARE section begins with the keyword DECLARE and ends when the keyword BEGIN signals the arrival of the EXECUTABLE section. You can declare types, constants, variables, exceptions, and cursors in any order, as long as they are declared before they are referenced in another definition. You declare subprograms last. A semi-colon terminates each definition.

Datatypes

PL/SQL provides a number of predefined datatypes for variables and constants. It also enables you to define your own types, which are subtypes of the predefined types. The types fall into the following three categories:

- Scalar. These include all string, number, and binary types. All of the SQL datatypes, which are the datatypes that you can store in the database, fall

into this category. To find out about these datatypes, see "Datatypes" in the *Oracle7 Server SQL Reference.*

- Composite. These are structured datatypes, which is to say data structures that have components you can address independently. The PL/SQL composite types are TABLE (which is distinct from both database and HTML tables) and RECORD. These types are explained later in this appendix.
- Reference. There is one kind of reference datatype—REF CURSOR—which is a pointer to a cursor. Cursors are explained later in this appendix. For more information on the REF CURSOR datatype, see "Using Cursor Variables" in the *PL/SQL User's Guide and Reference.*

For a list and explanation of all PL/SQL datatypes, see "Datatypes" in the *PL/SQL User's Guide and Reference.*

In many cases, you can convert from one datatype to another, either explicitly or automatically. The possible conversions and the procedure involved are explained in the *PL/SQL User's Guide and Reference* under "Datatype Conversion".

You can also define a variable so that it inherits its datatype from a database column or from another variable or constant, as explained in the next section.

Declaring Variables

For variables, provide the name, datatype, and any desired attributes, as follows:

```
cnum INTEGER(5) NOT NULL;
```

This declares a five-digit integer called *cnum* that will not accept nulls. The use of case above serves to distinguish keywords from identifiers; PL/SQL is not case-sensitive. NOT NULL is the only SQL constraint that you can use as a PL/SQL attribute.

Note: PL/SQL initializes all variables to null. Therefore, a NOT NULL variable, such as the above, produces an error if referenced before it is assigned a value.

Optionally, you can assign an initial value to the variable when you declare it by following the datatype specification with an assignment, as follows:

```
cnum INTEGER(5) := 254;
```

This sets *cnum* to the initial value of 254. Alternatively, you can use the keyword DEFAULT in place of the assignment operator := to achieve the same effect. For more information on setting defaults, see "Declarations" in the *PL/SQL User's Guide and Reference.*

Inheriting Datatypes

To have the variable inherit the datatype of a database column or of another variable, use the %TYPE attribute in place of a declared datatype, as follows:

```
snum cnum%TYPE;
```

This means that *snum* inherits the datatype of *cnum*. You can inherit datatypes from database columns in the same way, by using the notation *tablename.columnname* in place of the variable name. Normally, you do this if the variable in question is to place values in or retrieve them from the column. The advantages are that you need not know the exact datatype the column uses and that you need not change your code if the datatype of that column changes. If you do not own the table containing the column, precede the *tablename* with the *schemaname*, as described under *Ownership and Naming Conventions*. For more information on %TYPE assignments, see "Declarations" in the *PL/SQL User's Guide and Reference*.

Declaring Constants

You declare constants the same way as variables, except for the addition of the keyword CONSTANT and the mandatory assignment of a value. Constants do not take attributes other than the value. An example follows:

```
interest CONSTANT REAL(5,2) := 759.32;
```

Defining Types

User-defined types in PL/SQL are subtypes of existing datatypes. They provide you with the ability to rename types and to constrain them by specifying for your subtype lengths, maximum lengths, scales, or precisions, as appropriate to the standard datatype on which the subtype is based. For more information on the datatype parameters, see "Datatypes" in Chapter 2 of the *Oracle7 Server SQL Reference*. For more information on PL/SQL datatypes, see "Datatypes" in the *PL/SQL User's Guide and Reference*. You can also use the %TYPE attribute in defining a subtype. Here is an example:

```
SUBTYPE shortnum IS INTEGER(3);
```

This defines SHORTNUM as a 3-digit version of INTEGER. For more information see "User-Defined Subtypes" in the *PL/SQL User's Guide and Reference*.

Scope and Visibility

Nested subprograms, defined in the DECLARE section, can be called from either of the other sections, but only from within the same block where they are defined or within blocks contained in that block. Variables, constants, types, and

subprograms defined within a block are local to that block, and their definitions are not meaningful outside of it. Objects that are local to a block may be used by subprograms contained at any level of nesting in that same block. Such objects are global to the block that calls them.

The area of a program within which an object can be used is called the object's scope. An object's scope is distinct from its visibility. The former is the area of the program that can reference the object; the latter is the, generally smaller, portion that can reference it without qualification.

Qualification is used to override the default resolution of ambiguous references. An ambiguous reference can arise because objects or subprograms contained in different blocks can have the same names, even if they have overlapping scopes. When this happens, the reference by default means the object most local in scope—in other words, the first one PL/SQL finds by starting in the current block and working out to the enclosing ones. Qualification is the method used to override this. It is similar to the system of qualification used for database objects, as explained under *Ownership and Naming Conventions*. To qualify an object's name, precede it with the name of the subprogram where it is declared, followed by a dot, as follows:

```
relocate.transmit(245, destination);
```

This invokes a procedure called *transmit* declared in some subprogram called *relocate*. The subprogram *relocate* must be global to the block from which it is called.

Data Structures

PL/SQL provides two structured datatypes: TABLE and RECORD. It also provides a data structure called a cursor that holds the results of queries. Cursors are different from the other two in that you declare variables and constants to be of type TABLE or RECORD just as you would any other datatype. Cursors, on the other hand, have their own syntax and their own operations. Explanations of these types follow:

PL/SQL Tables

These are somewhat similar to database tables, except that they always consist of two columns: a column of values and a primary key. This also makes them similar to one-dimensional arrays, with the primary key functioning as the array index. Like SQL tables, PL/SQL tables have no fixed allocation of rows, but grow dynamically. One of their main uses is to enable you to pass entire columns of values as parameters to subprograms. With a set of such parameters, you can pass an entire table. The primary key is always of type BINARY_INTEGER, and the values can be of any scalar type.

You declare objects of type TABLE in two stages:

1. You declare a subtype using the following syntax:

```
TYPE type_name IS TABLE OF
     datatype_spec
     [ NOT NULL ]
     INDEX BY BINARY INTEGER;
```

   Where datatype_spec means the following:

```
datatype | variablename%TYPE | tablename.columnname%TYPE
```

   In other words, you can either specify the type of values directly or use the %TYPE attribute (explained under *Declaring Variables*) to inherit the datatype from an existing variable or database column.

2. You assign objects to this subtype in the usual way. You cannot assign initial values to tables, so the first reference to the table in the EXECUTABLE section must provide it at least one value.

When you reference PL/SQL tables, you use an array-like syntax of the form:

```
column_value(primary_key_value)
```

In other words, the third row (value) of a table called "Employees" would be referenced as follows:

```
Employees(3)
```

You can use these as ordinary expressions. For example, to assign a value to a table row, use the following syntax:

```
Employees(3) := 'Marsha';
```

For more information, see "PL/SQL Tables" in the *PL/SQL User's Guide and Reference*.

Records

As in many languages, these are data structures that contain one or more fields. Each record of a given type contains the same group of fields with different values. Each field has a datatype, which can be RECORD. In other words, you can nest records, creating data structures of arbitrary complexity. As with tables, you declare records by first declaring a subtype, using the following syntax:

```
TYPE record_type IS RECORD
     (fieldname datatype[, fieldname datatype]...);
```

The second line of the above indicates a parenthesized, comma-separated, list of fieldnames followed by datatype specifications. The datatype specifications can be direct or be inherited using the %TYPE attribute, as shown for TABLE and as explained under *Declaring Variables*.

You can also define a record type that automatically mirrors the structure of a database table or of a cursor, so that each record of the type corresponds to a row, and each field in the record corresponds to a column. To do this, use the %ROWTYPE attribute with a table or cursor name in the same way you would the %TYPE attribute with a variable,or column. The fields of the record inherit the column names and datatypes from the cursor or table. For more information, see "Records" and "%ROWTYPE Attribute" in the *PL/SQL User's Guide and Reference.*

Cursors

A cursor is a data structure that holds the results of a query (a SELECT statement) for processing by other statements. Since the output of any query has the structure of a table, you can think of a cursor as a temporary table whose content is the output of the query.

When you declare a cursor, you associate it with the desired query. When you want to use that cursor, you open it, executing the associated query and filling the cursor with its results. You then fetch each row of the query's output in turn for processing by other statements in the program. You can also use a cursor to update a table's contents. To do this, use a FOR UPDATE clause to lock the rows in the table. See "Using FOR UPDATE" in the *PL/SQL User's Guide and Reference* for more information. Sometimes, you may need to use cursor variables, which are not associated with a query until runtime. This is a form of dynamic SQL.

For more information on cursor variables, see "Using Dynamic SQL" in the *Oracle7 Server Application Developers Guide* and "Cursor Variables" in the *PL/SQL User's Guide and Reference.*

For more information on cursors in general, see "Cursors" in the *PL/SQL User's Guide and Reference.* See also "DECLARE CURSOR," "OPEN", and "FETCH" in the *Oracle7 Server SQL Reference.*

You can simplify some cursor operations by using cursor FOR loops. For more information on these, see "Using Cursor FOR Loops" in the *PL/SQL User's Guide and Reference.*

Exceptions

You also use the DECLARE section to define your own error conditions, called "exceptions". Explanation of this is deferred until the "EXCEPTION Section" portion of this appendix.

You must place all subprogram declarations at the end of the declare section, following all variable, constant, type, and exception declarations for the block. The syntax is as follows:

```
PROCEDURE procedure_name (param_name datatype, param_name datatype...)
     IS
     {local declarations}
     BEGIN {executable code}
     EXCEPTION
     END;
```

Note: For subprograms, the keyword DECLARE is omitted before the local declarations. Place local declarations before the keyword BEGIN, as shown.

The names you give the parameters in the declaration are the names that the procedure itself uses to refer to them. These are called the *formal parameters*. When the procedure is invoked, different variables or constants may be used to pass values to or from the formal parameters; these are called the *actual parameters*.

When calling the procedure, you can use each parameter for input of a value to the procedure, output of a value from it, or both. These correspond to the three *parameter modes*: IN, OUT, and IN/OUT. For more information, see "Parameter Modes" in the *PL/SQL User's Guide and Reference*.

When you call the procedure, you can match the actual to the formal parameters either implicitly, by passing them in the same order they are given in the declaration, or explicitly, by naming the formal followed by the actual parameter as shown:

```
transmit(destination => address);
```

This invokes a procedure called *transmit*, assigning the value of *address* as the actual parameter for the formal parameter *destination*. This implies that the parameter *destination* is used within the transmit procedure and that the parameter *address* is used outside of it. Usually, it is good programming practice to use different names for matching formal and actual parameters. For more information on this, see "Positional and Named Notation" in the *PL/SQL User's Guide and Reference*.

Functions are the same, except for the addition of a return value, specified as follows:

```
FUNCTION function_name (param_name, param_name datatype...)
     RETURN datatype IS
     {local declarations}
     BEGIN {executable code}
     EXCEPTION {local exception handlers}
     END;
```

Again, line breaks are only for readability. A RETURN statement in the executable section actually determines what the return value is. This consists of the keyword RETURN followed by an expression. When the function executes the RETURN statement, it terminates and passes the value of that expression to whichever statement called it in the containing block.

You can also use the RETURN statement without an expression in a procedure to force the procedure to exit.

For more information on procedures and functions, see "Declaring Subprograms" in the *PL/SQL User's Guide and Reference.*

The EXECUTABLE Section

The executable section is the main body of code. It consists primarily of SQL statements, flow control statements, and assignments. SQL statements are explained earlier in this appendix; assignments and flow-control statements are explained in the sections that follow.

Assignments

The assignment operator is :=. For example, the following statement assigns the value 45 to the variable a:

```
a := 45;
```

Character strings should be set off with single quotes (') as in all expressions. An example follows:

```
FNAME := 'Clair';
```

There are other examples of assignments in other parts of this appendix.

Flow Control

PL/SQL supports the following kinds of flow-control statements:

- IF statements. These execute a group of one or more statements based on whether a condition is TRUE.
- Basic loops. These repeatedly execute a group of one or more statements until an EXIT statement is reached.
- FOR loops. These repeatedly execute a group of one or more statements a given number of times or until an EXIT statement is reached.

- WHILE loops. These repeatedly execute a group of one or more statements until a particular condition is met or an EXIT statement is reached.
- GOTO statements. These pass execution directly to another point in the code, exiting loops and enclosing blocks as necessary. Use these sparsely, as they make code difficult to read and debug.

If you know other programming languages, you probably are familiar with most or all of these types of statements. The following sections describe the PL/SQL versions of them in greater detail. For more information on any of these, see "Control Structures" in the *PL/SQL User's Guide and Reference.*

You can nest flow control statements within one another to any level of complexity.

IF Statements

These are similar to the IF statement in many other languages, except that they use predicates, which are three-valued Boolean expressions like the SQL predicates discussed earlier in this appendix. In most respects, a Boolean NULL behaves like a Boolean FALSE, except that negation does not make it positive, but leaves it NULL.

The IF statement has the following forms:

```
IF <condition> THEN <statement-list>;
     END IF;
```

If the condition following IF is TRUE, PL/SQL executes the statements in the list following THEN. A semicolon terminates this list. END IF (not ENDIF) is mandatory and terminates the entire IF statement. Here is an example:

```
IF balance > 500 THEN send_bill(customer);
END IF;
```

We are assuming that send_bill is a procedure taking a single parameter.

```
IF <condition> THEN <statement-list>;
     ELSE <statement-list>;
     END IF;
```

This is the same as the preceding statement, except that, if that condition is FALSE or NULL, PL/SQL executes the statement list following ELSE instead of that following THEN.

```
IF <condition> THEN <statement-list>;
     ELSIF <condition> THEN <statement-list>;
     ELSIF <condition> THEN <statement-list>;.....
     ELSE <statement-list>;
END IF;
```

You can include any number of ELSIF (not ELSEIF) conditions. Each is tested only if the IF condition and all preceding ELSIF conditions are FALSE or NULL. As soon as PL/SQL finds an IF or ELSIF condition that is TRUE, it executes the associated THEN statement list and skips ahead to END IF. The ELSE clause is optional, but, if included, must come last. It is executed if all preceding IF and ELSIF conditions are FALSE or NULL.

NULL Statements

If you do not want an action to be taken for a given condition, you can use the NULL statement, which is not to be confused with database nulls, Boolean NULLs, or the SQL predicate IS NULL. The syntax of this statement is simply:

```
NULL;
```

The statement performs no action, but fulfills the syntax requirement that a statement list must follow every THEN keyword. In some cases, you can also use it to increase the readability of your code. For more information on the NULL statement, see "NULL Statement" in the *PL/SQL User's Guide and Reference*.

Basic Loops

A basic loop is a loop that keeps repeating until an EXIT statement is reached. The EXIT statement must be within the loop itself. If no EXIT (or GOTO) statement ever executes, the loop is infinite. An example follows:

```
credit := 0;
LOOP
     IF c = 5 THEN EXIT;
     END IF;
     credit := credit + 1;
END LOOP;
```

This loop keeps incrementing credit until it reaches 5 and then exits. An alternative to placing an exit statement inside an IF statement is to use the EXIT-WHEN syntax, as follows:

```
EXIT WHEN credit = 5;
```

This is equivalent to the earlier IF statement.

Note: The EXIT statement cannot be the last statement in a PL/SQL block. If you want to exit a PL/SQL block before its normal end is reached, use the RETURN statement. For more information, see "RETURN Statement" in the *PL/SQL User's Guide and Reference*.

FOR Loops

A FOR loop, as in most languages, repeats a group of statements a given number of times. The following FOR loop is equivalent to the example used for basic loops, except that it also changes a variable called interest.

```
FOR credit IN 1..5 LOOP
     interest := interest * 1.2;
END LOOP;
```

The numbers used to specify the range (in this case, 1 and 5) can be variables, so you can let the number of iterations of the loop be determined at runtime if you wish.

WHILE Loops

A WHILE loop repeats a group of statements until a condition is met. Here is a WHILE loop that is the equivalent of the preceding example:

```
credit := 1;
WHILE credit <= 5 LOOP
     interest := interest * 1.2;
     credit := credit + 1;
END LOOP;
```

Unlike some languages, PL/SQL has no structure, such as REPEAT-UNTIL, that forces a LOOP to execute at least once. You can create this effect, however, using either basic or WHILE loops and setting a variable to a value that will trigger the loop, as in the above example. For more information on loops, see "Iterative Control" in the *PL/SQL User's Guide and Reference.*

GOTO Statements

A GOTO statement immediately transfers execution to another point in the program. The point in the program where the statement is to arrive must be preceded by a label. A label is an identifier for a location in the code. It must be unique within its scope and must be enclosed in double angle brackets, as follows:

```
<<this_is_a_label>>
```

You only use the brackets at the target itself, not in the GOTO statement that references it, so a GOTO statement transferring execution to the above label would be:

```
GOTO this_is_a_label;
```

Note: An EXIT statement can also take a label, if that label indicates the beginning of a loop enclosing the EXIT statement. You can use this to exit several nested loops at once. See "Loop Labels" in the *PL/SQL User's Guide and Reference* for more information.

A GOTO statement is subject to the following restrictions:

- It must branch to an executable statement, not, for example, an END.
- It cannot branch to a point within the body of IF or a LOOP statement, unless it is contained in the body of that statement itself.
- It cannot branch to a subprogram or enclosing block of the present block (with one exception, explained shortly).
- It cannot branch from one IF statement clause to another. That is to say, it cannot jump between THEN, ELSIF, and ELSE clauses that are part of the same IF statement.
- It cannot branch from the EXCEPTION section to the EXECUTABLE section of the same block.
- It can, however, branch from the EXCEPTION section of a block to the EXECUTABLE section of an enclosing block, which is the exception to the third rule above.

The EXCEPTION Section

The EXCEPTION section follows the END that matches the BEGIN of the EXECUTABLE section and begins with the keyword EXCEPTION. It contains code that responds to runtime errors. An *exception* is a specific kind of runtime error. When that kind of error occurs, you say that the exception is *raised.* An *exception handler* is a body of code designed to handle a particular exception or group of exceptions. Exception handlers, like the rest of the code, are operative only once the code is compiled and therefore can do nothing about compilation errors.

There are two basic kinds of exceptions: predefined and user-defined. The predefined exceptions are provided by PL/SQL in a package called STANDARD. They correspond to various runtime problems that are known to arise often—for example, dividing by zero or running out of memory. These are listed in the *PL/SQL User's Guide and Reference* under "Predefined Exceptions".

The Oracle Server can distinguish between and track many more kinds of errors than the limited set that STANDARD predefines. Each of Oracle's hundreds of messages are identified with a number, and STANDARD has simply provided labels for a few of the common ones. You can deal with the other messages in either or both of two ways:

- You can define your own exception labels for specified Oracle messages using a pragma (a compiler directive). This procedure will be explained shortly.
- You can define a handler for the default exception OTHERS. Within that handler, you can identify the specific error by accessing the built-in functions SQLCODE and SQLERRM, which contain, respectively, the numeric code and a prose description of the message.

You can also define your own exceptions as will be shown. It is usually better, however, to use Oracle exceptions where possible, because then the conditions are tested automatically when each statement is executed, and an exception is raised if the error occurs.

Declaring Exceptions

PL/SQL predefined exceptions, of course, need not be declared. You declare user-defined exceptions or user-defined labels for Oracle messages in the DECLARE section, similarly to variables. An example follows:

```
customer_deceased EXCEPTION;
```

In other words, an identifier you choose followed by the keyword EXCEPTION. Notice that all this declaration has done is provide a name. The program still has no idea when this exception should be raised. In fact, there is at this point no way of telling if this is to be a user-defined exception or simply a label for an Oracle message.

Labeling Oracle Messages

If a previously-declared exception is to be a label for an Oracle error, you must define it as such with a second statement in the DECLARE section, as follows:

```
PRAGMA EXCEPTION_INIT (exception_name, Oracle_error_number);
```

A PRAGMA is a instruction for the compiler, and EXCEPTION_INIT is the type of PRAGMA. This tells the compiler to associate the given exception name with the given Oracle error number. This is the same number to which SQLCODE is set when the error occurs. The advantage of this over defining your own error condition is that you pass the responsibility for determining when the error has occurred and raising the exception to Oracle. You can find the numeric codes and explanations for Oracle messages in *Oracle7 Server Messages*.

User-Defined Exceptions

If the declared condition is not to be a label for an Oracle error, but a user-defined error, you do not need to put another statement referring to it in the DECLARE section. In the EXECUTABLE section, however, you must test the situation you

intend the exception to handle whenever appropriate and raise the condition manually, if needed. Here is an example:

```
IF cnum < 0 THEN RAISE customer_deceased;
```

You can also use the RAISE statement to force the raising of predefined exceptions. For more information, see "Error Handling" in the *PL/SQL User's Guide and Reference.*

Handling Exceptions

Once an exception is raised, whether explicitly with a RAISE statement or automatically by Oracle, execution passes to the EXCEPTION section of the block, where the various exception handlers reside. If a handler for the raised exception is not found in the current block, enclosing blocks are searched until one is found. If PL/SQL finds an OTHERS handler in any block, execution passes to that handler. An OTHERS handler must be the last handler in its block. If no handler for an exception is found, Oracle raises an unhandled exception error. Note: this does not automatically roll back (undo) changes made by the subprogram, which might leave the database in an undesirable intermediate state.

This is the syntax of an exception handler:

```
WHEN exception_condition THEN statement_list;
```

The exception is the identifier for the raised condition. If desired, you can specify multiple exceptions for the same handler, separated by the keyword OR. The exception can be either one the package STANDARD provided or one you declared. The statement list does what is appropriate to handle the error—writing information about it to a file, for example—and arranges to exit the block gracefully if possible. Although exceptions do not necessarily force program termination, they do force the program to exit the current block. You cannot override this with a GOTO statement. You can use a GOTO within an exception handler, but only if its destination is some enclosing block.

Note: If you have an error prone statement and want execution to continue following this statement, even when an exception occurs, put the statement, including the appropriate exception handlers, in its own block, so that the current block becomes the enclosing block.

Note: If an exception occurs in the DECLARE section or the EXCEPTION section itself, local exception handlers cannot address it; execution passes automatically to the EXCEPTION section of the enclosing block.

Storing Procedures and Functions in the Database

To have a procedure or function stored as a database object, you issue a CREATE PROCEDURE or a CREATE FUNCTION statement directly to the server using SQL*PLUS or Server Manager. The easy way to do this is to use your ordinary text editor to produce the CREATE statement and then to load it as a script. This process is explained under "Creating Stored Procedures and Functions" in the *Oracle7 Server Application Developers Guide*. This approach is recommended because you often create entire groups of procedures and functions together. These groups are called "packages" and are explained later in this appendix.

The syntax for these statements is slightly different than that used to declare subprograms in PL/SQL, as the following example shows:

```
CREATE PROCEDURE fire_employee (empno INTEGER) IS
    BEGIN
          DELETE FROM Employees WHERE enum = empno;
    END;
```

As you can see, the main difference is the addition of the keyword CREATE. You also have the option of replacing the keyword IS with AS, which does not affect the meaning. To replace an existing procedure of the same name with this procedure (as you frequently may need to do during development and testing), you can use CREATE OR REPLACE instead of simply CREATE. This destroys the old version, if any, without warning.

Privileges Required

A stored procedure or function (for the rest of this discussion, "procedure" shall mean "procedure or function" unless otherwise indicated or clear from context) is a database object like a table. It resides in a schema, and its use is controlled by privileges. To create a procedure and have it compile successfully, you must meet the following conditions:

- If the procedure is to be in your own schema, you must have the CREATE PROCEDURE or the CREATE ANY PROCEDURE system privilege. These privileges apply as well to functions.
- If the procedure is to be in a schema you do not own, you must have the CREATE ANY PROCEDURE system privilege.
- You must have the object privileges necessary to perform all operations contained in the procedure. You must have these privileges as a user, not through roles. If your privileges change after you have created the procedure, the procedure may no longer be executable.

To enable others to use the procedure, grant them the EXECUTE privilege on it using the SQL statement GRANT (see "GRANT" in Chapter 4 of the *Oracle7 Server SQL Reference*). When these users execute the procedure, they do so under your privileges, not their own. Therefore, you do not have to grant them the privileges to perform these actions outside the control of the procedure, which is a useful security feature. To enable all users to use the procedure, grant EXECUTE to PUBLIC. The following example permits all users to execute a procedure called show_product.

```
GRANT EXECUTE ON show_product TO PUBLIC;
```

Of course, the public normally does not execute such a procedure directly. This statement enables you to use the procedure in your PL/SQL code that is to be publicly executable. If multiple users access the same procedure simultaneously, each gets his own instance. This means that the setting of variables and other activities by different users do not affect one another.

For more information on privileges and roles, see "GRANT" in Chapter 4 of the *Oracle7 Server SQL Reference*. There are three versions of GRANT listed—one each for object privileges, system privileges, and roles.

For more information on storing procedures and functions in the database, see "Storing Procedures and Functions" in the *Oracle7 Server Application Developers Guide* and see "CREATE FUNCTION" and "CREATE PROCEDURE" in the *Oracle7 Server SQL Reference*.

Packages

A package is a group of related PL/SQL objects (variables, constants, types, and cursors) and subprograms that is stored in the database as a unit. Being a database object, a package resides in a schema, and its use is controlled by privileges. Among its differences from regular PL/SQL programs are that a package as such does not do anything. It is a collection of subprograms and objects, at least some of which are accessible to applications outside of it. It is the subprograms in the package that contain the executable code. A package has the following two parts:

- The package specification is the public interface to the package. It declares all objects and subprograms that are to be accessible from outside the package. Packages do not take parameters, so these constitute the entire public interface.
- The package body is the internal portion of the package. It contains all objects and subprograms that are to be local to the package. It also contains

definitions of the public cursors and subprograms. The package specification declares but does not define these.

One of the advantages of using packages is that the package specification is independent of the body. You can change the body and, so long as it still matches the specification, no changes to other code are needed, nor will any other references become invalid.

Packages cannot be nested, but they can call one another's public subprograms and reference one another's public objects.

Instantiation of Packages

It is important to realize that a package is instantiated once for a given user session. That is to say, the values of all variables and constants, as well as the contents and state of all cursors, in a package, once set, persist for the duration of the session, even if you exit the package. When you reenter the package, these objects retain the values and state they had before, unless they are explicitly reinitialized. Of course, another user has another session and therefore another set of values. Nonetheless, a global reinitialization of a package's objects for you does not take place until you disconnect from the database.

There is an exception, however. When one package calls another, execution of the second has a dependency on the first. If the first is invalidated, for example because its creator loses a privilege that the package requires, the second, while not necessarily invalidated, becomes deinstantiated. That is to say, all its objects are reinitialized.

Note: In PL/SQL, stored procedures and packages are automatically recompiled if changes to the database mandate it. For example, a change to the datatype of a column can automatically cascade to a variable referencing that column if the former is declared with the %TYPE attribute, but that change requires that the PL/SQL procedure declaring that variable be recompiled. So long as the PL/SQL code as written is still valid, the recompilation occurs automatically and invisibly to the user.

Creating Packages

To create a package, you use the SQL statement CREATE PACKAGE for the specification and CREATE PACKAGE BODY for the body. You must create the specification first. Sometimes, a package may consist of only public variables, types, and constants, in which case no body is necessary. Generally, however, you use both parts.

Note: Before you can create a package, the special user SYS must run the SQL script DBMSSTDX.SQL. The exact name and location of this script may vary

according to your operating system. Contact your database administrator if you are not sure this script has been run.

Creating the Package Specification

The syntax of the CREATE PACKAGE statement is as follows:

```
CREATE [OR REPLACE] PACKAGE package_name IS
    {PL/SQL declarations}
     END;
```

The optional OR REPLACE clause operates just as it does for stored procedure. The PL/SQL declarations are as outlined under *The DECLARE Section*, except that the keyword DECLARE is not used and that the subprogram and cursor declarations are incomplete. For subprograms, you provide only the name, parameters, and, in the case of functions, the datatype of the return value. For cursors, provide the name and a new item called the return type. This approach hides the implementation of these objects from the public while making the objects themselves accessible.

The syntax for declaring a cursor with a return type is as follows:

```
CURSOR c1 IS RETURN return_type;
```

The return type is always some sort of record type that provides a description of the cursor's output. The structure of this record is to mirror the structure of the cursor's rows. You can specify it using any of the following:

- A record subtype previously defined and in scope. For more information, see *Records*.
- A type inherited from such a record subtype using the %TYPE attribute. For more information, see *Declaring Variables*.
- A type inherited from a table, most likely the table the cursor queries, using the %ROWTYPE attribute. For more information, see *Records*.
- A type inherited from a cursor using the %ROWTYPE attribute. For more information, see *Records*.

For more information, see CREATE PACKAGE in Chapter 4 of the *Oracle7 Server SQL Reference*, "Packages" in the *PL/SQL User's Guide and Reference*, and "Using Procedures and Packages" in the *Oracle7 Server Application Developers Guide*.

Creating the Package Body

To create the package body, use the CREATE PACKAGE BODY statement. The syntax is as follows:

```
CREATE [OR REPLACE] PACKAGE BODY package_name IS
    {PL/SQL declarations}
```

```
END;
```

Since a package as such does not do anything, the PL/SQL code still consists only of a DECLARE section with the keyword DECLARE omitted. It is the subprograms within the package that contain the executable code. Variables, constants, types, and cursors declared directly (in other words, not within a subprogram) in the declare section have a global scope within the package body. Variables, constants, and types already declared in the package specification are public and should not be declared again here.

Public cursors and subprograms, however, must be declared again here, as their declarations in the specification is incomplete. This time the declarations must include the PL/SQL code (in the case of subprograms) or the query (in the case of cursors) that is to be executed. For subprograms, the parameter list must match that given in the package specification word for word (except for differences in white space). This means, for example, that you cannot specify a datatype directly in the specification and use the %TYPE attribute to specify it in the body.

You can create an initialization section at the end of the package body. This is a body of executable code—chiefly assignments—enclosed with the keywords BEGIN and END. Use this to initialize constants and variables that are global to the package, since otherwise they could be initialized only within subprograms, and you have no control of the order in which subprograms are called by outside applications. This initialization is performed only once per session.

For more information, see CREATE PACKAGE BODY in the *Oracle7 Server SQL Reference*, "Packages" in the *PL/SQL User's Guide and Reference*, and "Using Procedures and Packages" in the *Oracle7 Server Application Developers Guide*.

Overloading Subprograms

Within a package, subprogram names need not be unique, even at the same level of scope. There can be multiple like-named subprograms in the same declare section, provided that the parameters that they take differ in number, order, or datatype and that, when the procedures are called, the values passed by the calling procedure (the actual parameters) match or can be automatically converted to the datatypes specified in the declaration (the formal parameters). To find out which datatypes PL/SQL can convert automatically, look under "Datatype Conversion" in the *PL/SQL User's Guide and Reference*.

The reason this is permitted is so you can overload subprograms. Overloading permits you to have several versions of a procedure that are conceptually similar but behave differently with different parameters. This is one of the properties of

object-oriented programming. For more information on overloading, see "Overloading" in the *PL/SQL User's Guide and Reference*.

Database Triggers

Triggers are blocks of PL/SQL code that execute automatically in response to events. Database triggers reside in the database and respond to changes in the data. They are not to be confused with application triggers, which reside in applications and are beyond the scope of this discussion. Database triggers are a technology that for the most part has superseded application triggers.

You create triggers as you do stored procedures and packages, by using your text editor to write scripts that create them and then using SQL*Plus or Server Manager to run these scripts. A trigger is like a package in that:

- It takes no parameters as such. It refers to, responds to, and possibly affects the data in the database.
- It cannot be directly called like a procedure. To fire (execute) a trigger, you must make the database change to which it responds. If you only want to test the trigger, you can rollback (undo) the database change that you made after the trigger fires.

Triggers can be classified in three ways:

- INSERT triggers, UPDATE triggers, and DELETE triggers. This is a classification based on the statement to which the trigger responds. The categories are not mutually exclusive, meaning one trigger can respond to any or all of these statements.
- Row triggers and statement triggers. Any of the above statements can affect any number of rows in a table at once. A row trigger is fired once for each row affected. A statement trigger is fired once for each statement, however many rows it affects.
- BEFORE triggers and AFTER triggers. This specifies whether the trigger is fired before or after the data modification occurs.

As you can see, all three of these classifications apply to all triggers, so that there are, for example, BEFORE DELETE OR INSERT statement triggers and AFTER UPDATE row triggers.

Creating Triggers

The syntax of the CREATE TRIGGER statement is as follows:

```
CREATE [OR REPLACE] TRIGGER trigger_name
     BEFORE | AFTER
     DELETE | INSERT | UPDATE [OF column_list]
     ON table_name
     [ FOR EACH ROW [ WHEN predicate ] ]
     {PL/SQL block};
```

In the above, square brackets ([ ]) enclose optional elements. Vertical bars (|) indicate that what precedes may be replaced by what follows.

In other words, you must specify the following:

- A trigger name. This is used to alter or drop the trigger. The trigger name must be unique within the schema.
- BEFORE or AFTER. This specifies whether this is a BEFORE or AFTER trigger.
- INSERT, UPDATE, or DELETE. This specifies the type of statement that fires the trigger. If it is UPDATE, you optionally can specify a list of one or more columns, and only updates to those columns fire the trigger. In such a list, separate the column names with commas and spaces. You may specify this clause more than once for triggers that are to respond to multiple statements; if you do, separate the occurrences with the keyword OR surrounded by white space.
- ON table_name. This identifies the table with which the trigger is associated.
- PL/SQL Block. This is an anonymous PL/SQL block containing the code the trigger executes.

You optionally can specify the following:

- OR REPLACE. This has the usual effect.
- FOR EACH ROW [WHEN predicate]. This identifies the trigger as a row trigger. If omitted, the trigger is a statement trigger. Even if this clause is included, the WHEN clause remains optional. The WHEN clause contains a SQL (not a PL/SQL) predicate that is tested against each row the triggering statement alters. If the values in that row make the predicate TRUE, the trigger is fired; else it is not. If the WHEN clause is omitted, the trigger is fired for each altered row.

Here is an example:

```
CREATE TRIGGER give_bonus
     AFTER UPDATE OF sales
     ON salespeople
     FOR EACH ROW WHEN sales > 8000.00
     BEGIN
     UPDATE salescommissions SET bonus = bonus + 150.00;
     END;
```

This creates a row trigger called give_bonus. Every time the sales column of the salespeople table is updated, the trigger checks to see if it is over 8000.00. If so, it executes the PL/SQL block, consisting in this case of a single SQL statement that increments the bonus column in the salescommissions table by 150.00.

Privileges Required

To create a trigger in your own schema, you must have the CREATE TRIGGER system privilege and one of the following must be true:

- You own the table associated with the trigger.
- You have the ALTER privilege on the table associated with the trigger.
- You have the ALTER ANY TABLE system privilege.

To create a trigger in another user's schema, you must have the CREATE ANY TRIGGER system privilege. To create such a trigger, you precede the trigger name in the CREATE TRIGGER statement with the name of the schema wherein it will reside, using the conventional dot notation.

Referring to Altered and Unaltered States

You can use the correlation variables OLD and NEW in the PL/SQL block to refer to values in the table before and after the triggering statement had its effect. Simply precede the column names with these variables using the dot notation.

If these names are not suitable, you can define others using the REFERENCING clause of the CREATE TRIGGER statement, which is omitted from the syntax diagram above for the sake of simplicity. For more information on this clause, see CREATE TRIGGER in the *Oracle7 Server SQL Reference*.

Note: if a trigger raises an unhandled exception, its execution fails and the statement that triggered it is rolled back if necessary. This enables you to use triggers to define complex constraints. If the effects of the trigger have caused a change in the value of package body variables, however, this change is not reversed. You should try to design your packages to spot this eventuality. For more information, see "Using Database Triggers" in the *Oracle7 Server Application Developers Guide*.

Enabling and Disabling Triggers

Just because a trigger exists does not mean it is in effect. If the trigger is disabled, it does not fire. By default, all triggers are enabled when created, but you can disable a trigger using the ALTER TRIGGER statement. To do this, the trigger must be in your schema, or you must have the ALTER ANY TRIGGER system privilege. Here is the syntax:

```
ALTER TRIGGER trigger_name DISABLE;
```

Later you can enable the trigger again by issuing the same statement with ENABLE in place of DISABLE. The ALTER TRIGGER statement does not alter the trigger in any other way. To do that you must replace the trigger with a new version using CREATE OR REPLACE TRIGGER. For more information on enabling triggers, see ALTER TRIGGER in the *Oracle7 Server SQL Reference.*

For more information on triggers generally, see "Using Database Triggers" in the *Oracle7 Server Application Developer's Guide* and CREATE TRIGGER and DROP TRIGGER in the *Oracle7 Server SQL Reference.*

ORCL000999

APPENDIX

C Introduction To HTML

This appendix discusses the basic concepts of HyperText Markup Language, or HTML. This appendix is an introduction to HTML and provides important information on how to set up, format, and define HTML documents.

This appendix covers the following topics:

- *What is HTML?*
- *Getting Started*
- *Document Structure*
- *Body Tags*
- *List Tags*
- *Hypertext Linking*
- *Reviewing Changes to Your HTML Document*
- *Adding Style to Your HTML Document*
- *Special HTML Tags*
- *Tables*
- *Forms*
- *Creating Your Own HTML Document*
- *More Information about HTML*

Note: This appendix contains examples of formatted elements, such as underlined text. Some of these examples may not appear correctly on all on-line viewing systems.

What is HTML?

Hypertext Markup Language (HTML) is the standard language used for creating hypermedia documents on the Web. HTML documents can be viewed by many different Web browsers, of varying abilities, in a simple, portable way. When a document is coded in HTML, a browser can interpret the HTML to identify the elements of the document and to render it. The use of HTML allows documents to be formatted for presentation using fonts and line justification appropriate for the system on which it is displayed.

Most documents have common elements, such as a title, paragraphs, or lists. Using HTML *tags* you can label these elements as you are writing. HTML tags provide the browser with a minimum of presentation information, while keeping the integrity of information in the document. All the reader needs is a formatting tool, a Web browser, which interprets the HTML tags and produces an on-screen display that approximates the intent of the document creator.

With most methods of documentation, the writer of a document has strict control over the look and feel of a document. With HTML, the *reader* (subject to the capabilities of the Web browser) has control over the look and feel of a document. HTML allows you to mark titles or paragraphs with HTML instructions or tags, and then leaves the interpretation of these tags up to the browser. For example, one browser may indent the beginning of each paragraph, and another may leave only a blank line. The user of a particular browser may also have some control over the specific fonts used.

HTML tags can be divided into two main categories:

- tags that define how the body of the document is to be displayed by the browser
- tags that define information about the document such as the title

Remember, the power of HTML is that your document can be viewed on a variety of browsers, on most platforms, and formatted to suit any reader.

How Are HTML Documents Created?

HTML documents can be created using any text editor (emacs, textedit, or vi on UNIX machines; DOS and Macintosh machines have a variety of simple text editors or HTML-specific programs). Choose the editor you are most comfortable with to write your HTML document.

For example, HTML editors such as SoftQuad's "HoTMetaL" allow the creation of HTML documents graphically in "what you see is what you get" (WYSIWYG) mode. In addition, many traditional word processing packages have add-ons or integrated HTML output capabilities.

To create dynamic pages that retrieve information from an Oracle7 Server, you can generate HTML using the PL/SQL utility packages provided by the Developer's Toolkit. See "Appendix B: The PL/SQL Developer's Toolkit," for more information.

Getting Started

All HTML tags begin with a < (left angle bracket), and end with a > (right angle bracket). There is usually a beginning tag and an ending tag.

An example is the title tag which surrounds the text that is designated as the document's title:

```
<TITLE>All the Hockey Greats</TITLE>
```

Tags are usually paired as follows:

```
<TITLE> and </TITLE>
```

The ending tag looks like the beginning tag except that a forward slash precedes the text within the bracket. In this example, the tag <TITLE> tells the Web browser to use a title format, and the </TITLE> tells the browser that the title heading is complete.

A few tags, such as <P>, which is a paragraph delimiter, do not need an end tag, but most do.

HTML is not case sensitive; therefore, the previous tags could look like this:

```
<title>All the Hockey Greats</title>
```

The convention used in this document is to capitalize all HTML format tags.

Note: Extra spaces, tabs, or returns that you have added by hand are highly discouraged, and will be lost. HTML only interprets tabs, extra spaces, and

returns enclosed by the <PRE> </PRE> tags. See "Preformatted Tag" for more information on this HTML option.

Document Structure

When a browser receives a document, it determines how it should be interpreted. The very first tag you need in your HTML document is the <HTML> structure tag. This declares that the content of your document is written with HTML. A minimal HTML document would look like this:

```
<HTML>...the content of the document...
</HTML>
```

Head Tag

The head tag can be used right after the HTML declaration, or not at all in your document. This tag represents the prologue to the rest of the file. Avoid putting any text into the document <HEAD> tag. This tag is placed immediately before and after the <TITLE> tag, as shown in the following example:

```
<HTML>
<HEAD>
<TITLE>All the Hockey Greats</TITLE>
</HEAD>
```

Note: Technically, the start and end tags for <HTML>, <HEAD>, and <BODY> are not needed. However, they are recommended because the head and body structure tags allow a browser to determine certain properties of a document, such as the <TITLE>, without having to parse, or go through the whole document.

Title Tag

Most browsers display the contents of the <TITLE> tag in the title bar of the window containing the document, and in the bookmark file of the browser if it supports one. The title, surrounded by <TITLE> and </TITLE> tags, is placed between the head tags, as shown above. The title of a document does not appear in the contents of the document window, however. You must separately indicate it as a heading inside the body of the document if you want it to appear there.

Body Tags

The body tags specifically identify the body components in an HTML document. The body of an HTML document can contain links, text, and formatting information inside of the <BODY> and </BODY> tags.

Body Tag

The body of the document should be marked off with the <BODY> and </BODY> tags. This is the part of the document that is displayed as the page of text and graphics on your Web browser.

Heading Levels

When writing an HTML document, organize the text by heading levels to reflect its structure and organization. The first heading would be level 1, the next sub heading level 2, and so on. Most browsers recognize up to six heading levels, with six distinct styles. Heading levels above 6 are indistinguishable from one another.

The largest heading is a level 1 heading. The syntax of the head 1 is as follows:

```
<H1>Hockey Greats on Offense</H1>
```

Other headings can be created as follows:

```
<Hx>Text here</Hx>
```

where *x* is a number between 1 and 6 specifying the heading level. For example, if your next heading level is a level 3, the syntax would look like the following:

```
<H3>Hockey Defense </H3>
```

Paragraph Tag

Unlike most word processors, HTML usually ignores carriage returns. Word wrapping can occur at any point in your source file. Therefore, paragraphs must be separated with the <P> tag. If you do not separate your paragraphs with the <P> tag, your document will look like one long paragraph.

Preformatted Tag

The preformatted tag, <PRE>, allows you to present text formatted specifically to a screen. The preformatted text ends at the closing </PRE> tag. Within the preformatted text:

- Line breaks move to the next line
- The <P> tag is moved to the next line

- Horizontal tabs move in multiples of eight
- A fixed width font is used for all characters after the <PRE> tag.

Avoid using tags that define paragraph formatting, such as headings or address, within the <PRE> tags. They will have no effect.

Let's incorporate some of the previous examples to show what the document looks like with a title, a couple of heading level tags, and a few paragraphs:

```
<HTML>
<HEAD>
<TITLE>All the Hockey Greats</TITLE>
</HEAD>
<BODY>
<H1>All the Hockey Greats Before 1970</H1>
<H2>The Original Six Teams</H2>
This section deals with all of the hockey legends before the expansion.<P>
The game was very different for these players.<P>
There is no way New Jersey would have won the Stanley Cup. New Jersey just
would not have had the talent to do it.<P>
Chicago would still be on top.<P>
All would be well.<P>
</BODY>
</HTML>
```

The result would display something like this:

All The Hockey Greats Before 1970

The Original Six Teams

This section deals with all of the hockey legends before the expansion.

The game was very different for these players.

There is no way New Jersey would have won the Stanley Cup. New Jersey just would not have had the talent to do it.

Chicago would still be on top.

All would be well.

Note: The title, "All the Hockey Greats," would not show within the document itself. On most browsers it would be displayed in the title bar.

Forced Line Breaks

The <BR> tag forces a line to break. The best example of the use of this tag is for formatting addresses, or some other sequence of lines where you don't want the browser to add extra spacing. For example:

```
Sandy's Super Sundaes<BR>
123 Main Street<BR>
Anytown, USA<BR>
```

BlockQuote

The <BLOCKQUOTE> tag is used to contain text quoted from another source. The quote will be indented approximately 8 spaces. For example:

```
My favorite hockey saying is<P>
<BLOCKQUOTE>
Today is a great day for hockey.
</BLOCKQUOTE>
But I'm not sure if Bob Johnson really said it like that.</P>
```

This would appear something like the following:

My favorite hockey saying is

> Today is a great day for hockey.

But I'm not sure if Bob Johnson really said it like that.

Summary of Basic HTML Tags

The following table lists the basic HTML tags and their corresponding values:

Opening	*Closing*	*Definition*
<HTML>	</HTML>	An entire HTML document
<HEAD>	</HEAD>	The prologue of the document
<TITLE>	</TITLE>	Title of the document
<BODY>	</BODY>	Content of the document
<H1>	</H2>	First level heading
<H2>	</H2>	Second level heading
<H3>	</H3>	Third level heading
<H4>	</H4>	Fourth level heading
<H5>	</H5>	Fifth level heading
<H6>	</H6>	Sixth level heading
<P>		Paragraph
<PRE>	</PRE>	Preformatted text
 		Forced line break
<BLOCKQUOTE>	</BLOCKQUOTE>	Text quoted from another source

The previous section provides all you need to know to get started with HTML. At this point you can write a simple document in HTML. However, the following sections show you how to enhance your HTML pages to present information in many different fashions.

List Tags

There are three basic lists in HTML:

ordered

These lists have numbered items.

unordered

These lists have bullets to mark each item.

definition

These lists alternate a term with its definition

You can create nested lists with indents using the ordered or unordered tags. Simply place a second list (complete with its own start and end tags) within the first lists enclosing tags. Whether the nested list uses the same markers (numbers or bullets,) depends on the browser; some track the number of nests you use and change the markers of each successive nesting to blocks or other symbols. See the following example in the section "Nested Lists".

Ordered Lists

In an ordered list, the browser automatically inserts numbers. Therefore, if you insert or delete an item in your ordered list, the numbers will reflect the change automatically.

An ordered list begins with <OL> and ends with </OL>. The individual list items are started with the <LI> tag. The following is an example of an ordered list:

```
<OL>
<LI>Gordie Howe
<LI>Rocket Richard
<LI>Howie Morenz
</OL>
```

Unordered Lists

In an unordered list the browser typically uses bullets or dashes to indicate the items in your list. (Each browser has its own way of indicating an unordered list)

An unordered list begins with <UL> and ends with </UL>. The following is an example of an unordered list:

```
<UL>
<LI>Gordie Howe
<LI>Rocket Richard
<LI>Howie Morenz
</UL>
```

Nested Lists

Here is an example of how to nest a list within another:

```
<HTML>
<HEAD>
<TITLE>All the Hockey Greats</TITLE>
</HEAD>
<BODY>
<H1>Hockey Greats before Expansion</H1>
<H2>The Original Six Teams</H2>
This section deals with all of the hockey legends before the expansion.<P>
<UL>
<LI>Gordie Howe
<LI>Rocket Richard
<LI>Howie Morenz
<UL>
<LI>great player
<LI>good stickhandler
</UL>
<LI>Bobby Orr
</UL>
</BODY>
</HTML>
```

Here's what the example would look like:

Hockey Greats before Expansion

The Original Six Teams

This section deals with all of the hockey legends before the expansion.'

- Gordie Howe
- Rocket Richard
- Howie Morenz
  - great player
  - good stickhandler
- Bobby Orr

Note: When you create nested lists using HTML, you do not need to indent the nested HTML components. You may wish to do so for clarity, however.

Definition Lists

The definition list tags <DL>...<D/DL> enclose both the defined term (identified with the <DT> tag), and the definition of that term (identified with the <DD> tag). Most browsers format the definition on a separate line from the term. The following is an example of a definition list:

```
<DL>
<DT>Slapshot:
<DD>A shot used to drill the goalie at speeds up to 100 mph.
<DT>Wristshot:
<DD>A shot used to scare the goalie after the slapshot.
</DL>
```

The output looks like this:

Slapshot:
A shot used to drill the goalie at speeds up to 100 mph.

Wristshot:
A shot used to scare the goalie after the slapshot.

Hypertext Linking

Hypertext linking is the key characteristic that makes the Web appealing to users. By adding hypertext links, called *anchors* in your HTML document, you can create a highly intuitive information flow and guide the user directly to the information he or she needs.

Anchors have a standard format that allows any Web browser to interpret a link and perform the proper function (called a method) for that type of link. Links can refer to other documents, specific locations within the same document, or can perform operations, such as retrieving a file using FTP for display by the browser. URLs can refer to a specific location by an absolute pathname, or can be relative to the current document, which is more convenient when managing a large site.

Note: You can use hypertext links to navigate through a document or to move from document to document. However, HTML does not support returning you to the anchor point of a link within a document. If you use a hypertext link within a document, and then use the Back button, you do not return to the anchor, but to the previous point that you reached through a link.

What is a URL?

HTML uses what are called Uniform Resource Locators (URLs) to represent hypermedia links and links to network services within documents. The first part of the URL before the colon specifies the access method. The part of the URL after the colon is interpreted specifically according to the access method. In general two forward slashes after the colon indicate a machine name.

The general format of a URL is:

```
method://machine-name/path/foo.html
```

The following example would fetch the document `index.html` from the server `www.acme.com` using the HTTP protocol.

```
http://www.acme.com/index.html
```

A Uniform Resource Locator (URL) has the following format:

method://servername:port/pathname#anchor

The components of the URL are as follows:

method

is the name of the operation that is performed to interpret this URL. The most common methods are:

file

Read a file from the local disk. The filename is interpreted on the *user*'s machine, so this can be used to display any file that resides on the user's disk.

For example: `file:/home/jjones/jjones.html` displays the file `jjones.html` from the directory `/home/jjones` on the local machine.

http

Access a page over the network by way of the HTTP protocol. (This is the most common method, usually used to get an HTML document.)

For example: `http://www.acme.com/`
accesses Acme's home page.

mailto

Activate a mail session to the specified username and host.

For example: `mailto:jjones@us.acme.com` mails a message to jjones if the browser supports mail creation. Note that the `mailto` method does not require double forward slashes after the colon.

ftp

Retrieve a file using anonymous FTP from a server.

For example, `ftp://hostname/directory/filename`

servername

is optional and indicates the full hostname of a machine connected to the network. For example, `www.oracle.com` is the fully qualified hostname of Oracle's web server. If a servername is not specified, the URL is a *relative* link, and it is assumed that the file is on the same server that was used to display the current page. An IP address may be used instead of a hostname, although it is not recommended to build content with embedded IP addresses.

port

is the TCP port number that the web server is running on. The default is 80 if `port` is not specified. This parameter is not used in most URLs.

pathname

is the relative or absolute pathname of the document being accessed by this URL. Web servers can be configured to interpret certain pathnames differently. For instance, CGI applications work by configuring the HTTP server to recognize that files within certain directories should be executed instead of being returned to the browser.

For example: `http://www.acme.com/index.html`

In this example, an HTTP connection is made to `index.html`, which is the name of the file to be accessed on the server `www.acme.com` using port 80 (the default). The file could have a full UNIX-style pathname to indicate a document contained in a lower level directory. If there is no directory component in the pathname, the document must be located in the server's document root directory which is configured by the server administrator. If the pathname part of the URL is missing, many servers provide a directory listing of the document root directory or access a specific 'top level' file (usually `index.html`).

#anchor

The named anchor points to a specific location within an HTML page. In addition to specifying a document name, specifying #anchor will cause most browsers to move the top of the display to the point referred to by the anchor. These anchor names are inserted into documents with the NAME tag as explained under "Linking to Sections on Pages" later in this appendix.

Structure of an Anchor Link within a Document

So far, we've discussed what a URL looks like. To cause a link to be displayed for the user to access, an anchor link must be embedded in the document text. The HTML syntax which allows this is:

```
<A HREF="URL">text_to_be_displayed_highlighted</A>
```

The <A HREF="*URL*"> tag opens the anchor link and the </A> tag closes it. All the text between those tags is displayed highlighted in some way by a web browser. A common technique is to display it underlined and in blue or some other user-selected contrasting color.

The URL part of the tag refers to the text of the URL reference as defined in the previous section. The URL text does not appear on the user's screen; it is only used when the user activates the link, usually by clicking on it with the mouse.

An example HTML segment:

```
For interesting products see
<A HREF="http://www.acme.com">Acme's home page.</A>
```

This line would produce on the user's screen:

For interesting products see **Acme's home page**.

Graphics within HTML Documents

One of the most compelling features of the Web is the ability to embed references to graphics and other data types within a document using the <IMG...ISMAP> tags. This adds a very lively character to your pages and makes them visually interesting.

There are two ways to use graphics from within an HTML document. The first is by embedding them within the document itself, so the user's screen will display the graphics within the context of the other elements of your document (such as explanatory text). This is the most common technique used by HTML designers and is called an "inline image". The syntax for specifying this is:

```
<IMG SRC="URL" ALT="text"
   ALIGN=["top"|"middle"|"bottom"|"texttop"] ISMAP>
```

The elements in this syntax statement are as follows:

URL

is the same syntax as any other URL, as explained above. This is the way the browser accesses the actual image data file, which should be in a format

supported by the browser. Currently GIF and JPEG formats are supported by most browsers. Specifying the URL is required.

ALT="text"

will cause the string *text* to be displayed if the browser is incapable of displaying images, or if image display is turned off. This is a way of 'labelling' the image the user would be seeing if image display was turned on. ALT is an optional keyword; if it is missing, no text will be displayed if images are turned off on the browser. Most browsers put some sort of icon on the screen to indicate an image would normally be there. Using the ALT tag is recommended so that your page is compatible with text-only browsers such as Lynx.

ALIGN

is used with one of the keywords to tell the browser where to place the next block of text. This allows a certain amount of creativity in the layout of your page. If this is not specified, most browsers put the image on the left side of the screen and fill in following text to the right of it.

ISMAP

tells the server that this image is a bitmap and allows the user to click on a location on the image to cause a URL to be accessed directly. Image maps are an advanced HTML feature and require server side configuration to function.

For example, the following line will cause the file `logo.gif` to be fetched from server `www.oracle.com` and the text `Oracle Logo` to be displayed if the user has graphic display turned off.

```
<IMG SRC="http://www.oracle.com/logo.gif" ALT="Oracle Logo">
```

Linking to Sections on Pages

You can link to a different area or section of your document by using a hidden reference marker to that specific section. This provides a quick way to move through sections of the document without having to scroll up or down. Once you click on that link, the hidden reference places you in the section, and the browser presents the hidden marker line as the first line on the screen.

To create a link, follow these steps:

1. Create a named anchor marker in each section title that you might want to jump to. The HTML syntax is as follows:

```
<A NAME="named_anchor"> Text_to_link_to</A>
```

2. In the following example, a hidden reference marker is placed in a head2

level:

```
<H2><A NAME="intro">Introduction to Hockey</A></H2>
```

3. Create the link by entering:

```
<A HREF="#named_anchor">Text</A>
```

4. For example, the link might appear as:

```
<A HREF="#intro">Gordie Howe</A>
```

The "#" symbol instructs your web browser to look through the HTML document for a named anchor called "intro".

When the user clicks on "Gordie Howe," the browser displays the heading "Introduction to Hockey" at the top of the screen.

Note: As specified earlier in the URL syntax, a link to a section can appear in the same document or in another document. The example here describes a link to a specific section of the same document.

Reviewing Changes to Your HTML Document

At some point, you will want to view the changes you have made to your HTML document. The following steps show you how to view your changes:

1. After you have edited your HTML file in your favorite text editor, save the file.
2. If you are currently displaying the document in a browser, you will need to reload it to see the changes. Select File -->Reload on the browser window. (In some browsers, reload will be under another menu item.)

The browser will read in the new file information and display the file with the changes you have made.

Adding Style to Your HTML Document

HTML tags offer several text styles so that you can emphasize any text in your document. The following is a short list of the most often used tag styles:

- bold
- italics
- underline

- mono spaced (typewriter style)

You can also use combinations of styles (for example, bold and italics)

Style	Element or Tag	Result
Bold	<B>I want this text bold</B>	**I want this text bold**
Italics	<I>I want this text italics</I>	*I want this text italics*
Underlined	<U>I want this underlined</U>	I want this underlined
Mono spaced(typewriter)	<TT>I want this text typed</TT>	`I want this text typed`

You can add style to text that appears anywhere in the HTML document. Combinations of styles can also be used, provided you assign all the closing tags needed.

```
<I>Hockey</I> is <B>Life</B>.<P>
```

becomes a paragraph that looks like this:

Hockey is **Life**.

The style tags surround the words they affect, in conjunction with other tags such as headings.

Be careful how many style tags you use on one page. If you use too many, the text may become difficult to read.

Special HTML Tags

The following tags or characters are options provided to make your HTML document more robust.

Address Tag

The <ADDRESS> tag is used to specify the author of a particular HTML document, and a way to contact the author (such as an email address). The syntax is as follows:

```
<ADDRESS>
address_of_author
</ADDRESS>
```

Escape Sequences

The following ASCII characters have special meaning within HTML and cannot be used in regular text:

- left angle bracket <
- right angle bracket >
- ampersand &
- double quote "

To use these characters you must use these escape sequences:

For	Use
<	<
>	>
&	&amp
"	"

Table C-1:

There are several more escape sequences to support accented characters, such as the umlaut or the tilde. A full list of supported characters can be found in any number of reference manuals on HTML syntax. See also the list of on-line HTML references at the end of this appendix.

Tables

Tables in HTML organize data by row and column. Tables can contain a wide range of content, such as headers, lists, paragraphs, or figures. They can include any element or tag in HTML.

Cells can be merged across rows or columns.

Basic Table Tags

This section describes the basic table tags and their meanings.

Table: `<TABLE>...</TABLE>`

This is the main wrapper for all the other table tags. Other table tags can be ignored if they aren't wrapped inside of the <TABLE></TABLE> tags. By default, tables have no borders. Borders are added if the BORDER attribute is specified. See the next section "Basic Table Attributes".

Table Row: `<TR></TR>`

The number of rows in a table is specified by how many <TR> tags are contained within it. <TR> can have both the ALIGN and VALIGN attributes, which if specified become the default alignments for all cells in this row. See the next section, "Basic Table Attributes".

Table Data: `<TD></TD>`

This specifies a standard table data cell. Table data cells must only appear within table rows. Each row need not have the same number of cells specified, because short rows will be padded with blank cells on the right. A cell can contain any of the HTML tags normally present in the body of an HTML document. The default alignment of table data is ALIGN=left and VALIGN=middle. These alignments can be overridden by any alignments specified in the containing <TR> tag.

Note: Row alignments are overridden by any attributes assigned to a cell.

By default, lines inside of table cells can be broken up to fit within the overall cell width. Use the NOWRAP attribute described in the section, "Basic Table Attributes," to prevent line breaking for that cell.

Table Header: `<TH></TH>`

The table header cells are identical to data cells in all respects, except that header cells are in a bold font and have a default ALIGN=center.

Caption: `<CAPTION>... </CAPTION>`

This optional tag represents the caption for a table. The <CAPTION> tags should appear inside the <TABLE></TABLE> tags but not inside table rows or cells. The default alignment for the <CAPTION> tag is ALIGN=top but can be explicitly set to ALIGN=bottom. Like table cells, any document body HTML tags can appear in a caption. Captions are always horizontally centered with respect to the table, and they may have their lines broken to fit within the width of the table.

ORCL001017

Basic Table Attributes

This section lists the basic table attributes and their meanings.

BORDER: This attribute appears in the Table tag. If present, borders are drawn around all table cells. If absent, there are no borders. By default space is left for borders, so a table has the same width with or without the border attribute.

ALIGN: If the ALIGN attribute appears inside a <CAPTION> </CAPTION> tag, it controls whether the caption appears above or below the table. It can have the values top or bottom. The default attribute is ALIGN=top.

When appearing inside a <TR>, <TH>, or <TD> tag, ALIGN controls whether text inside the table cell(s) is aligned to the left side of the cell, the right side of the cell, or centered within the cell. Values are left, center, and right.

VALIGN: The VALIGN attribute appears inside a <TR>, <TH>, or <TD> tag. This attribute controls whether text inside the table cell(s) is aligned to the top of the cell, the bottom of the cell, or vertically centered within the cell. It can also specify that all the cells in the row should be vertically aligned to the same baseline. Values are top, middle, bottom, and baseline.

NOWRAP: If the NOWRAP attribute appears in any table cell, the lines within this cell cannot be broken to fit the width of the cell. Be cautious in use of this attribute as it can result in excessively wide cells.

COLSPAN: The COLSPAN attribute can appear in any table cell and specifies how many columns of the table a specified cell should span. The default COLSPAN for any cell is 1.

ROWSPAN: The ROWSPAN attribute can appear in any table cell and specifies how many rows of the table this cell should span. The default ROWSPAN for any cell is 1. A span that extends into rows that were not specified with a <TR> tag will be truncated.

COLSPEC: The COLSPEC attribute can be used when needed to exert control over column widths, either by setting explicit widths or by specifying relative widths. Specify the table width explicitly or as a fraction of the current margins.

Example Table

The following is a table using some of the table tags and attributes we have previously discussed:

```
<TABLE BORDER>
<CAPTION ALIGN=bottom> Table #1 </CAPTION>
<TR><TD ROWSPAN=2> </TD> <TH COLSPAN=2>Average</TH></TR>
```

```
<TR><TH>Height</TH><TH>Weight</TH></TR>
<TR><TD>Males</TD><TD ALIGN=center> 69 </TD>
<TD ALIGN=center>150 </TD></TR>
<TR><TD>Females</TD><TD ALIGN=center> 64 </TD>
<TD ALIGN=center>130 </TD></TR>
</TABLE>
```

The table will look like this:

	Average	
	Height	**Weight**
Males	69	150
Females	64	130

	Average	
	Height	**Weight**
Males	69	150
Females	64	130

Table #1

```
<TABLE BORDER>
<CAPTION ALIGN=bottom> Table #1 </CAPTION>
<TR><TD ROWSPAN=2> </TD> <TH COLSPAN=2>Average</TH></TR>
<TR><TH>Height</TH><TH>Weight</TH></TR>
<TR><TD>Males</TD><TD ALIGN=center> 69 </TD>
<TD ALIGN=center>150 </TD></TR>
<TR><TD>Females</TD><TD ALIGN=center> 64 </TD>
<TD ALIGN=center>130 </TD></TR>
</TABLE>
```

Forms

HTML pages can be formatted in any fashion, but remain read-only. The HTML *form* feature brings the added advantage of being interactive. An HTML form lets the Web user enter comments and specify database search criteria.

When a form is interpreted by a Web browser, a special graphical user interface (GUI) screen is created with text entry fields, buttons, checkboxes, pull-down menus, and scrolling lists. When the Web user fills out the form and presses a button indicating the form should be *submitted*, the information on the form is sent to an HTTP server for processing by a CGI (Common Gateway Interface) program.

When you write a form, each of your input items has an <INPUT> tag. When the user places data in these items in the form, that information is encoded into the form data and is known as the "value".

All form elements have Name and Value attributes. Data are sent as NAME=VALUE pairs, separated by ampersands (&), where name is given in the NAME attribute, and value is given in the VALUE attribute, or replaced by the user.

This section illustrates the basic use of HTML forms. Forms can be used for simple table searches or complex queries to relational databases.

Forms Syntax

All forms begin with <FORM> and end with </FORM>.

The syntax is as follows:

```
<FORM METHOD="get|post" ACTION="URL">Form_elements_and_other_HTML
</FORM>
```

METHOD

The request method supplies the data to the program. There are two request methods that can be used to access your forms. Depending on which request method you use, you will receive the encoded results of the form in a different way.

- GET: Information from a form is appended onto the end of the URL being requested. Your CGI program receives the encoded form input in the environment variable QUERY_STRING. Use of the GET method is discouraged.
- POST: This request method transmits all form input information immediately after the requested URL. Your CGI program will receive the encoded form input on standard input. The server will *not* send you an EOF on the end of the data; instead use the environment variable CONTENT_LENGTH to determine how much data you should read from standard input. This is the preferred method.

ACTION

ACTION specifies the URL being requested from the form. This URL will almost always point to a CGI script to decode the form results. If you are referring to a script on the same server as the form, you can use a relative URL.

TEXTAREA

The <TEXTAREA> tag is used to allow a user to enter more than one line of text in a form. The following is an example of the TEXTAREA tag:

```
<TEXTAREA NAME="address" ROWS=14 COLS=60>
Chicago Blackhawks
1800 Madison Ave
Chicago, Il 60612
</TEXTAREA>
```

The attributes included within the <TEXTAREA> tag are used to initialize the field's value. The </TEXTAREA> tag is always required even if the field is initially blank.

The following are attributes of <TEXTAREA> and determine the visible dimensions of the field in characters:

- NAME-user-defined name
- ROWS-height in characters of TEXTAREA
- COLS-width in characters of TEXTAREA

If you want text to appear within the text area, enter it between the start and end <TEXTAREA>tags.

INPUT

The <INPUT> tag allows you to input a single word or line of text, with a default width of 20 characters. It is usually preceded with some descriptive text.

The following are attributes of <INPUT>:

- CHECKED-When present indicates that a checkbox or radio button is selected.
- MAXLENGTH-The maximum number of characters that will be accepted as input. This limits the number of characters a user can type into the field. The form will give an error beep if the user tries to enter too many characters. This can be greater that specified by SIZE, in which case the field will scroll appropriately. The default is unlimited.
- NAME-Symbolic field name used when transferring the form's contents. This attribute is always needed and should uniquely identify this field.
- SIZE-Specifies how large an area to allocate, in characters, on the screen.
- SRC-A URL specifying an image for use only with TYPE=IMAGE.



- **TYPE**-Defines the type of data the field accepts. Defaults to free text. Several types of fields can be defined with the TYPE attribute:

  CHECKBOX

  Used for simple Boolean attributes, or for attributes that can take multiple values at the same time. The latter is represented by a number of checkbox fields each of which has the same NAME. Each occurrence of a checkbox is either ON or OFF. For a multi-valued attribute, there is a checkbox for each attribute, indicating whether the attribute applies. The various attributes are associated with one another by the fact that each attribute checkbox uses the same name. Each selected checkbox generates a separate NAME=VALUE pair in the submitted data, even if this results in duplicate names. The default value for CHECKBOX is ON.

  HIDDEN

  No field is presented to the user, but the content of the field is sent with the submitted form. This value may be used to transmit state information about client-server interaction, or for password information.

  IMAGE

  An image field upon which you can click with a pointing device, causing the form to be immediately submitted. The coordinates of the selected point are measured in pixel units from the upper left corner of the image, and are returned (along with the other contents of the form) in two NAME=VALUE pairs. The x-coordinate is submitted under the name of the field with .x appended, and the y-coordinate is submitted under the name of the field with .y appended. Any value attribute is ignored. The image itself is specified by the SRC attribute, exactly as for the <IMG> tag.

  PASSWORD

  Same as TEXT attribute, except that password text is not displayed as it is entered.

  RADIO

  For attributes which can take a single value from a set of alternatives. Each radio button field in the group should be given the same NAME. Only the selected radio button in the group generates a NAME=VALUE pair in the submitted data. Radio buttons require an explicit VALUE attribute.

  RESET

  This is a button that when pressed resets the form's fields to their specified initial values.

SUBMIT

This button, when pressed, submits the form. You can use the VALUE attribute to provide a non-editable label to be displayed on the button. The default label is application-specific. The NAME attribute, if used, passes a name=value pair along with the submitted form.

TEXT

Single line text entry fields. Use in conjunction with the SIZE and MAXLENGTH attributes. Use the TEXTAREA tag for text fields that can accept multiple lines.

- VALUE-Assigns an initial default value for the field, or the value when checked for checkboxes and radio buttons. This attribute is required for radio buttons.

Form Selection Menus

There are three types of selection menu tag for forms:

- **Select**: user selects from a fixed set of values represented by the *option* tag. This tag is usually displayed with a pull down menu.
- **Select single**: same as Select, but display is presented with a window with three items displayed at once. If there are more than three options, the window will have a scroll bar.
- **Select multiple**: allows multiple items to be selected from the menu.

SELECT

The SELECT tag allows the user to select a value from a fixed list. This is usually presented as a pull down menu.

The SELECT tag has one or more options between the start<SELECT> and end</SELECT> tag. By default the first option is displayed in the menu. The following is an example of a <SELECT> tag:

```
<FORM>
<SELECT NAME=group>
<OPTION>Gretzky
<OPTION>Messier
<OPTION>Coffey
</SELECT>
</FORM>
```

SELECT SINGLE

The SELECT SINGLE tag is the same as the SELECT tag, but options are displayed in a window with three items shown at once. If there are more than three options, the window will have a scroll bar. The SIZE tag within the SELECT tag specifies how many options will be shown in the window. The following is an example of a <SELECT SINGLE> tag:

```
<FORM>
<SELECT SINGLE NAME=group SIZE=3>
<OPTION>Gretzky
<OPTION>Messier
<OPTION>Coffey
<OPTION>Kurri
</SELECT>
</FORM>
```

In this example, the first three names would appear in the window, and a scroll bar would scroll to the last name.

SELECT MULTIPLE

The SELECT MULTIPLE tag is the same as the SELECT SINGLE tag, but the user can select more than one option in the window. The SIZE tag specifies how many lines appear in the window, and the MULTIPLE tag specifies how many options can be selected.

The following is an example of SELECT MULTIPLE:

```
<FORM>
<SELECT MULTIPLE NAME=group SiZE=3MULTIPLE=2>
<OPTION>Gretzky
<OPTION>Messier
<OPTION>Coffey
<OPTION>Kurri
</SELECT>
</FORM>
```

Note: On some browsers, it may be necessary to hold down the CONTROL or SHIFT key to select multiple items.

If multiple items are selected, they each get passed to the server with the same name. The decoding script has to be able to recognize multiple values associated with the same name.

Creating Your Own HTML Document

Now that most of HTML has been demystified, you can create your own HTML documents. Use any browser to view the HTML document you have created.

The following is an example of an HTML document that you can type word for word, or modify as you wish:

```
<HTML>
<HEAD>
<TITLE>Chicago Blackhawks: A Love Story</TITLE>
</HEAD>
<BODY>Chicago Blackhawks: A Love Story
<P>This is a story about my beloved Blackhawks. Year after painful year
they amass teams that could potentially win a Stanley Cup. But year after
insidiously painful year, they manage to lose in the first or second round
of the playoffs. If you have any suggestions as to how to reconcile this
evil, many fans would be grateful. Still we live on to love the Blackhawks.
<H2>Some Players We Have Loved</H2>
<UL>
<LI><A HREF="Hull.html">Bobby Hull</A>
<LI><A HREF="Makita.html">Stan Makita</A>
<LI><A HREF="Espo.html">Tony Esposito</A>
<LI><A HREF="Jr.html">Jeremy Roenick</A>
</UL>
<H2>Some Players We Love to Hate</H2>
<OL>
<LI><A HREF="Clark.html">Wendall Clark</A>
<LI><A HREF="Gilmour.html">Doug Gilmour</A>
<LI><A HREF="Domi.html">Tie Domi</A>
<LI><A HREF="cheapshot.html">Ulf Samuellson</A>
</OL>
<H2>List of Love Letters to Our Team</H2>
<UL>
<LI><A HREF="http://www.love2hawks.com/">Love Letters</A>
<LI><A HREF="http://www.fromNHL.com/">Letters from NHL</A>
</UL>
<HR><A HREF="Write">Write</A>to me.
Click below to send me Comments.
<BR> <A HREF="mailto:oldstadium@madison.com">
<I>crazed4hawks, oldstadium@madison.com</I>
</A>
</BODY>
</HTML>
```

More Information about HTML

For online information on HTML, the following URL sites provide a wealth of information:

http://www.ncsa.uiuc.edu/demoweb/html-primer.html

http://union.ncsa.uiuc.edu/HyperNews/get/www/html.html

http://fire.clarkson.edu/doc/html/htut.html

http://ugweb.cs.ualberta.ca/~gerald/guild/html.html



Index

%ROWTYPE attribute, B-19
%TYPE attribute, B-16, B-18

A

aliases
  for database objects, B-8
  for database tables, B-10
ALTER TRIGGER statement (SQL), B-35
anchors
  structure of, C-13
  used in hypertext linking, C-10
applets
  defined, 1-15
  embedded in Web pages, 1-15
  wrapping in server-side Java, 4-15
application developers, 1-1
applications
  components of, 4-1
  specifying, 4-4
  structure of in Java, 4-18
arrays
  implemented as tables in PL/SQL, B-17
assignments
  initial, B-15
  values to parameters (PL/SQL), B-20
  values to variables (PL/SQL), B-21
authentication, 1-5
automatic recompilation (PL/SQL), B-30

B

basic authentication, 1-5
blocks (PL/SQL), B-14
Boolean logic
  IF statements and (PL/SQL), B-22
  Three-Valued in SQL, B-6
browsers
  Java-enabled, 1-15
bytecode, 1-15

C

caching
  files, 1-5
cattributes, 6-5
cattributes parameter
  use in passing exact text, 6-5

certifying authorities, 1-9
CGI, 1-10
  environment variables, 4-4, 4-6, 4-31, 6-57
  executing from LiveHTML, 4-32
  LiveHTML can call, 1-17
  variables used by WRB, 1-12
changes
  viewing in HTML document, C-15
Common Gateway Interface (CGI), 1-10
compiler directives (PL/SQL), B-26
compression, 1-6
Connect-String, 4-5
constants
  declaring (PL/SQL), B-16
constraints (SQL), B-7
cookies, 6-79
correlation variables
  in SQL statements, B-10
  in triggers, B-35
CREATE FUNCTION statement (SQL), B-28
CREATE PACKAGE BODY statement (SQL), B-31
CREATE PACKAGE statement (SQL), B-31
CREATE PROCEDURE privilege (SQL), B-28
CREATE PROCEDURE statement (SQL), B-28
CREATE TABLE statement (SQL), B-7
crippled LiveHTML, 4-30
cursor variables (PL/SQL), B-19
cursors (PL/SQL), B-19
  in package specifications, B-31

D

data structures
  in PL/SQL, B-17
databases
  access control, 4-5
  changing data content of, B-8
  connecting to, 4-5, 4-16
  connecting to in Java, 4-18
  constraining, B-7
  creating objects in using PL/SQL, B-13
  creating tables in, B-7
  database triggers, B-33
  interfacing to, B-5
  missing data in, B-6
  nulls used in, B-6
  referencing objects in, B-7
  removing data from, B-9
  retrieving data from, B-6
  storing code in, B-13, B-28, B-33
datatypes
  composite, B-15
  converting, B-15
  in OWA_PATTERN, 6-66
  inheritance of (PL/SQL), B-16
  OWA_TEXT uses, 6-75
  PL/SQL, B-14
  PL/SQL Agent limitations, 4-12
  PL/SQL vs. Java, 4-18
  PL/SQL Web Toolkit, 6-5
  pointer, B-15
  RECORD (PL/SQL), B-18
  scalar, B-14
  SQL, B-7
  structured, B-15
  TABLE (PL/SQL), B-17
  user-defined, B-16
  user-defined in PL/SQL, B-18
DATE_GMT, 4-31
DATE_LOCAL, 4-31
DBMS_SQL package, B-13
DBMSSTDX.SQL script, B-30
DCDs, 4-16
  format of, 4-5
  how specified, 1-12, 4-4
  optimizing for multiple, 6-3

declarations
constants (PL/SQL), B-16
subprograms (PL/SQL), B-20
types (PL/SQL), B-16
variable (PL/SQL), B-15
DECLARE section (PL/SQL), B-14
DELETE statement (SQL), B-9
digest authentication, 1-5
digital signatures, 1-9
DNS, 1-4
DOCUMENT_NAME, 4-31
DOCUMENT_URL, 4-31
domain names, 1-4
domains
restricting access to specified, 1-6
dot notation
for database objects, B-7
for PL/SQL objects and subprograms, B-16–B-17

E

encryption
communications, 1-7
password, 1-5
public key, 1-8
session keys, 1-8
environment variables, 4-31
retrieving, 6-57
error handling
Java, 4-18
LiveHTML, 4-30
Oracle7, B-25
PL/SQL, B-25
PL/SQL Agent, 4-13
triggers and, B-35
exceptions (PL/SQL), B-25
declared in DECLARE section, B-19
declaring, B-26
predefined, B-25
user-defined, B-26
EXECUTABLE section (PL/SQL), B-21
EXECUTE privilege (SQL), B-28
EXIT statement (PL/SQL), B-24
extensions
filename, 1-6
extra path information, 4-3

F

fields
in PL/SQL records, B-18
file sharing, 1-5
files
access control, 1-5
caching, 1-5
compression of, 1-6
filename extesions, 1-6
language of, 1-6
MIME types, 1-6
firewalls, 1-4
flow control (PL/SQL), B-21
FOR loops (PL/SQL), B-24
FOR UPDATE clause (PL/SQL), B-19
forced line breaks, C-6
foreign keys
joins using, B-9
Form tags, C-22
Input, C-22
TextArea, C-22
form tags, 6-42
Forms, C-20
selection menus, C-24
syntax, C-21
forms
multivalued parameters, 4-9
parameter passing through, 4-8
frame tags
HTML frame tags, 6-29
FTP, C-12
functions, B-20
grouped into packages, B-29
return values of (PL/SQL), B-20
storing in database, B-28

G

GET, 4-6–4-7
GOTO statement (PL/SQL), B-24
  error handling and, B-27
graphics, 6-78
  image maps, 1-16
  within HTML documents, C-13

H

host names, 1-4
HTML
  basic concepts, C-2
  case sensitivity, C-3
  comments in, 4-30
  dynamically generating, 4-4
  extending, 6-63
  form tags, 6-42
  introduction, C-1
  Java generates dynamic, 4-25
  list tags, 6-31
  print tags, 6-6
  structure tags, 6-7
  table tags, 6-51
HTML Tags
  ADDRESS, C-16
  BLOCKQUOTE, C-7
  BODY, C-5
  BOLD, C-15
  CAPTION, C-18
  definition list, C-10
  escape sequence, C-17
  getting started, C-3
  HEAD, C-4
  heading level, C-5
  I, C-15
  lists, C-8
  mono spaced, C-15
  nested list, C-9
  ordered list, C-8
  P, C-5
  PRE, C-5
  TABLE, C-18
  TD, C-18
  TH, C-18
  TITLE, C-4
  TR, C-18
  U, C-15
  unordered list, C-9
HTTP, C-11
HyperText Markup Language, C-1
HyperText Procedures
  body tags, 6-13
  character format tags, 6-37
  definition of, 6-1
  form tags, 6-42
  head-related tags, 6-10
  list tags, 6-31
  physical format tags, 6-41
  printing procedures, 6-6
  structural tags (HTML), 6-7
  table tags, 6-51

I

IF statement (PL/SQL), B-22



image maps
defined, 1-16
dynamic, 1-16, 6-78
initialization section (PL/SQL packages), B-32
input parameters, B-20
INSERT statement (SQL), B-8
int_array, 6-75
IP addresses, 1-3
restricting access to specified, 1-6
ISMAP attribute, C-14

J

Java
application structure, 4-18
client side, 1-16
database connections, 4-18
error handling, 4-18
generating HTML from, 4-25
NULL handling in, 4-19
overview of, 1-15
PL/SQL combined with, 1-16
PL/SQL datatypes in, 4-18
proscribes direct memory management, 1-15
server side, 1-16, 4-15
wrappers for PL/SQL, 4-16
joins, B-9
natural, B-9
outer, B-10

L

labels
PL/SQL, B-24
languages
file formats and, 1-6
specified for WRBXs, 1-12
LAST_MODIFIED, 4-31
links
to other documents, C-10
to other sections, C-14
list tags, 6-31
definition, C-10
nested, C-9
ordered, C-8
types of, C-8
unordered, C-9
Listener, 1-3
LiveHTML
environment variables, 4-31
executing programs from, 4-32
overview of, 1-16
PL/SQL Agent can be called from, 1-17
tags, 4-30
loops (PL/SQL), B-23
exiting nested, B-24
FOR, B-24
WHILE, B-24

M

mailto, C-11
memory
management of, 1-5
proscribed in Java, 1-15
multi_line, 6-75

N

naming conventions
PL/SQL, B-17
tables, B-18
SQL, B-7
natural joins, B-10
NLS
file formats and, 1-6
NOT NULL constraint (SQL), B-7
PL/SQL use of, B-15
NULL statement (PL/SQL), B-23

nulls
as Boolean values, B-6
database contains, B-6
differentiated from NULL statement, B-23
outer joins generate, B-10
prohibiting, B-7, B-15
SQL handling of, B-6
UPDATE statement (SQL) generates, B-9
variables initialized to (PL/SQL), B-15

O

objects
ownership of database, B-7
oracle.html(Java package), 4-15
oracle.plsql(Java package), 4-15
oracle.rdbms(Java package), 4-15
ORACLE_HOME, 4-5
outer joins, B-10
output parameters, B-20
overloading
PL/SQL procedures, 4-12
overloading subprograms
PL/SQL, B-32
owa, 6-3
OWA_COOKIE, 6-79
OWA_IMAGE, 6-78
OWA_PATTERN, 6-63
OWA_TEXT, 6-75
OWA_UTIL Package, 6-55
OWAINS.SQL, 6-3

P

packages, B-29
creating, B-30
specification, B-31
parameters
actual (PL/SQL), B-20
formal (PL/SQL), B-20
in PL/SQL, B-20
in PL/SQL packages, B-32
modes of (PL/SQL), B-20
multivalued, 4-9
passing through forms, 4-8
parent keys, B-4
PATH_INFO, 4-4, 4-6
PL/SQL
blocks, B-14
datatypes
Java encapsulation of, 4-18
declarations, B-14
DECLARE section, B-14
embedding dynamic output in static Web pages, 1-17
EXCEPTION section, B-25
EXECUTABLE section, B-21
flow control in, B-21
generating dynamically, B-13
GOTO restrictions, B-25
Java combined with, 1-16
overloading subprograms, B-32
parameter passing, 4-4, 4-7
procedures
overloading, 4-12
scope of objects and subprograms, B-16
sections of, B-14
subprograms, B-14
user-defined datatypes in, B-16
PL/SQL Agent, 4-4
connect to database, 4-16
datatype limitations, 4-12
environment variables, 4-6
error handling, 4-13
Java circumvents, 4-15
LiveHTML can call, 1-17
multiple DCDs, 6-3
PL/SQL Agents, 1-13

PL/SQL packages
bodies, B-31
Java wrappers for, 4-16
PL/SQL tables
datatype limitations, 4-12
PL/SQL Web Toolkit, 6-1
datatypes used, 6-5
installation, 6-3
installing, 6-3
ownership of, 6-3
procedures and functions, 6-1
security considerations, 6-4
pl2java, 4-16
ports, 1-3
specifying in URLs, 4-3
POST, 4-6–4-7
pragmas (PL/SQL), B-26
predicates, B-6
in IF statements (PL/SQL), B-22
primary keys
enforcing, B-7
print tags, 6-6
private synonyms, B-8
privileges
object, B-5
required to create triggers, B-35
stored procedures require, B-28
stored procedures use, B-28
system, B-5
PRIVUTIL.SQL, 6-4
procedures
declaring, B-20
getting started, 6-5
grouped into packages, B-29
storing in database, B-28
proxy servers, 1-4
public key encryption, 1-8
public synonyms, B-8
PUBUTIL.SQL, 6-4

Q

queries (SQL), B-6
multiple tables used in, B-9
stored in cursors (PL/SQL), B-19
within other queries (subqueries), B-11
query strings, 4-3
QUERY_FORM, 4-9
QUERY_STRING, 4-6
QUERY_STRING_UNESCAPED, 4-31

R

RAISE statement (PL/SQL), B-27
range variables, B-10
records
PL/SQL datatype, B-18
ref cursors (PL/SQL), B-15
regular expessions, 6-63
reload, C-15
REQUEST_METHOD, 4-6–4-7
restriction, 1-6
RETURN statement (PL/SQL), B-20
return values
cursor (PL/SQL packages), B-31
functions (PL/SQL), B-20
row_list, 6-76

S

schemas, B-7
scope
of PL/SQL objects and subprograms, B-16
PL/SQL packages and, B-32
SCRIPT_NAME, 4-4, 4-6



security
certifying authorities, 1-9
DCD, 4-5
digital signatures, 1-9
encryption, 1-7
file access, 1-5
PL/SQL Web Toolkit, 6-4
proxy Internet connections, 1-4
public key encryption, 1-8
session keys, 1-8
SSL, 1-7
SELECT statement (SQL), B-6
server extensions, 4-2
Server Side Includes
same as LiveHTML, 1-16
session keys, 1-8
SGML comments, 4-30
SHTML, 4-30
sockets
defined, 1-4
SQL, B-5
dynamic, B-13
predicates, B-6
standards, B-5
Three-Valued Logic (TRUE, FALSE, NULL) in, B-6
SQLCODE, B-25
SQLERRM, B-25
SSL
Overview of, 1-7
STANDARD package (PL/SQL), B-25
stored procedures, B-28
string matching, 6-63
structure tags, 6-7
subprograms (PL/SQL), B-14
declaring, B-20
error handling and, B-27
GOTO restrictions, B-25
overloading, B-32
parameters of, B-17
resolving ambiguous references, B-17
scope and visibility of, B-16
subqueries (SQL), B-11
subtypes (PL/SQL), B-16
using to create PL/SQL records, B-18
using to create PL/SQL tables, B-18
synonyms
for database objects, B-8

T

Table attributes
Align, C-19
Border, C-19
Colspan, C-19
Colspec, C-19
Nowrap, C-19
Rowspan, C-19
Valign, C-19
table tags, 6-51
Tables
in HTML, C-17
tables
database, B-7
aliases for, B-8, B-10
HTML, 4-9
ownership of, B-7
passing as parameters in PL/SQL, B-17
PL/SQL, B-18
PL/SQL datatype, B-17
stored in cursors (PL/SQL), B-19
TCP/IP, 1-3
text/x-server-parsed-html, 4-30
text-only browsers, C-14
Three-Valued Logic (TRUE, FALSE, NULL), B-6
addressing in Java, 4-19
triggers, B-33
altering, B-33
creating, B-33
enabling and disabling, B-35

U

Uniform Resource Locator (URL), C-11, C-13
unique keys, B-4
UPDATE statement (SQL), B-9
URL, C-13
URLs, C-10
  GET vs. POST, 4-7
  interpretation of, 4-3
  parameter ordering, 4-9
  specifying PL/SQL Agent with, 1-13
  specifying ports in, 4-3
  specifying secure connections with, 1-4
  use of "owa" in, 1-13
  use virtual file names, 1-4
  used in applications, 4-2
usernames
  specified in DCDs, 4-5
users
  schemas and, B-7

V

variables
  assigning values to (PL/SQL), B-21
  cursor, B-15, B-19
  declaring (PL/SQL), B-15
  in triggers, B-35
  initializing
    in PL/SQL packages, B-32
  instantiation in PL/SQL packages, B-30
vc_array, 6-75
virtual file systems
  defined, 1-4
Virtual Machines (VMs), 1-15
virtual paths, 4-3
visibility
  of PL/SQL objects and subprograms, B-16
  PL/SQL packages and, B-32

W

Web Listener, 1-3
Web Listeners
  control WRBX load, 1-12
  file type negotiation, 1-6
  interpret URLs, 1-2
  memory management, 1-5
  Overview of, 1-3
Web pages
  applications and, 4-2
  dynamic, 1-16, 4-2, 4-4
    from Java, 4-25
    including in LiveHTML, 4-30
  dynamic data embedded in static, 1-17
  embedded, 4-30
Web Request Broker (WRB), 1-11
WebServer
  executing Java on, 1-16
  Overview of, 1-2
WebServer Administrators, 1-1
WebServer Manager
  overview of, 1-9
WHILE loops (PL/SQL), B-24
WRB, 1-11
  CGI environment variables and, 4-6
  overview of, 1-2
WRB cartridges
  defined, 1-2
WRB Dispatchers
  how invoked, 4-4
WRB Services
  defined, 1-2
WRBXs
  allocation of requests to, 1-12
  connect to database, 4-16
  connecting to database, 1-13
  data passed to, 1-12
  defined, 1-2
  single-threaded, 1-12

Exhibit 48

Oracle WebServer 2.0

Technical Note
March 1996

ORACLE®
Enabling the Information Age™

Oracle WebServer 2.0

This Technical Note is one of a series of publications from Oracle Internet Server Products, a product group within the Server Technologies Division of Oracle Corporation. This document is written for web software developers and other audiences familiar with basic Web technology. Specifically, no introduction to the Internet or Web will be given here. By reading this document, you will gain a technical understanding of how Oracle WebServer 2.0 is designed, and how you can use it to write high performance applications for the web. Many acronyms and technical terms are used in this document. For your convenience, we have included a glossary in appendix A.

Overview

Oracle WebServer 2.0 lets you publish hypermedia documents and deploy media-rich applications on the World Wide Web and corporate intranets using the industry standard HTTP protocol. Unlike mere HTTP servers, Oracle WebServer 2.0 offers a fully integrated platform for application development, mass-storage, and system and site-management, all using fully integrated state-of-the-art relational database technology and established industry standards such as HTTP, SSL, and HTML. Oracle WebServer 2.0 offers an open, extensible application development platform for the Web, enabling software developers to safely and easily add high-speed, custom server extensions to take full advantage of any and all modern technologies such as Java, distributed objects and virtual reality to name a few, without having to resort to the low-level network interfaces employed by most other HTTP servers today.

An Application Platform for the Web

Oracle WebServer 2.0 provides an open platform for developing web applications. For customers with heavy investments in the Oracle7 RDBMS, we offer the integrated PL/SQL Agent, an extension to Oracle WebServer 2.0, which allows you to develop complete web applications using Oracle7 stored procedures. For customers who wish to use Java, we offer a fully integrated Java Interpreter, enabling you to develop complete web applications using server-side Java classes. The Java Interpreter includes classes for native Oracle7 access via PL/SQL and dynamic creation of HTML documents. For customers who wish to employ their own programming language of choice, Oracle WebServer 2.0 offers an easy-to-use API allowing you to build your own

server-side plugins ranging from existing legacy applications written in COBOL to modern CORBA or COM/OLE based object-oriented systems written in for instance C++ or Smalltalk. Oracle WebServer 2.0 allows developers to easily integrate their back-end applications using our open API to the Oracle Web Request Broker.

Architecture

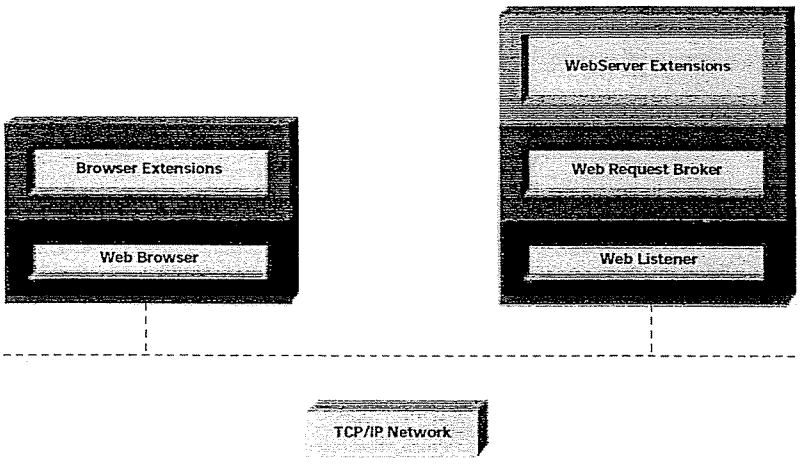
From a web browsers point of view, Oracle WebServer 2.0 looks just like any other web server. The industry standard HTTP 1.0 protocol with optional SSL based data-encryption is used to establish a short-term connection and request an object from the server. When the object is received by the client, the MIME type of the object determines how it is handled: either processed natively by the client, or handed off to a helper application or client-side plugin (extension) for further processing. This defines a complete HTTP transaction and is most often (but certainly not limited to) a request for and delivery of an HTML page, which is displayed in the browsers main viewing area.

Underneath the covers of the Oracle WebServer's HTTP server component, the Oracle Web Listener, interesting things happen to fulfill a request from a web browser. The following sections provides an in-depth description of the three major components of Oracle WebServer 2.0: the *Web Listener, Web Request Broker and the WebServer SDK.*

Oracle Web Listener

The Web Listener is responsible for delivering incoming requests to the Web Request Broker, where they are processed or dispatched to services built with the WebServer SDK.

Oracle WebServer 2.0 Architecture

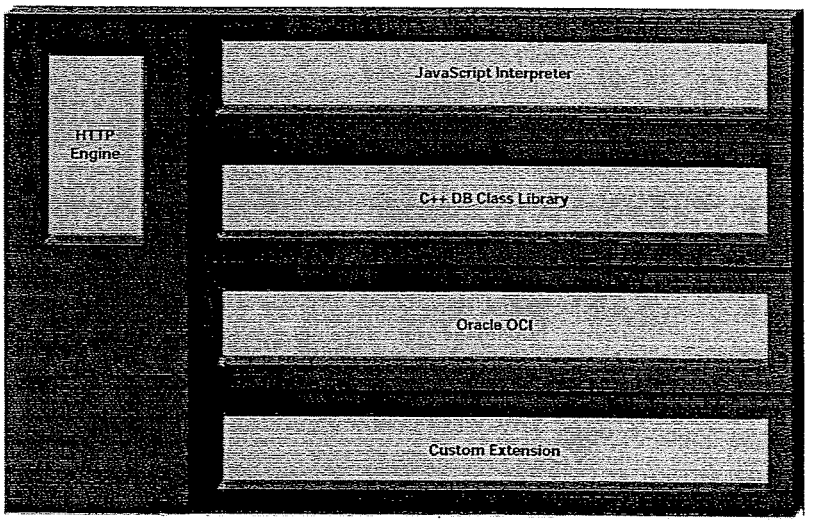


ORCL000763

Oracle's HTTP Server is designed to be a very scalable, high-performance network daemon. It is implemented as a multi-threaded, single-process asynchronous engine, allowing multiple concurrent requests to be processed at the same time using standard HTTP or HTTP over SSL. The HTTP Server is the network layer component of the Oracle WebServer. It listens for incoming HTTP requests, delivers static files and runs simple CGI programs, and hands off everything else to the Web Request Broker (WRB), where a request is handled by a server extension. The first generation of websites on the Internet simply offered static HTML files to their visitors. To enable dynamic content, web designers are mainly using CGI to extend the capabilities of their web server. Although easy to use, CGI has some serious limitations, primarily the high overhead of spawning a new process for each incoming request, which causes severe degradation in scalability and throughput.

Today, several webservers have defined an API for implementing extensions to the basic HTTP engine. The typical HTTP API enables a custom extension to be linked into the HTTP server process and allows it to intercept and fulfill a client request. Examples of such APIs are Netscape NSAPI, Spyglass ADI, and Microsoft IS-API. These APIs are becoming increasingly popular because they address the serious shortcomings of the original CGI protocol.

The problem with the HTTP API approach is that all extensions are linked into the same executable and thus share the same memory, process, and signal handling space. Extensions are free to clobber each other, and a poorly designed extension can bring the entire system down. Furthermore, technical support immediately becomes a nightmare: who are you going to call when you have linked together components from 5 different vendors ? What if the extension from vendor A is multi-threaded, and the extension from vendor B isn't thread-safe ?



Typical HTTP server extended through API

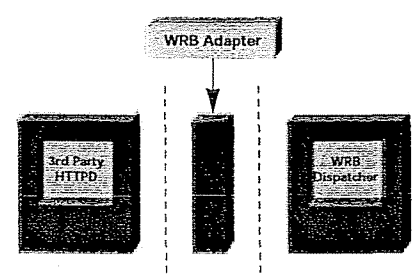
Oracle's answer to these issues is a new and improved architecture for the server, which effectively layers Oracle WebServer 2.0 into two fundamental components: a HTTP engine (the Oracle Web Listener) and a high-speed dispatch mechanism (the Web Request Broker) which routes requests to server extensions running in separate processes.

Being the industry leader in server and networking technology, Oracle fully understands how to optimize performance and throughput in an environment like the World Wide Web. Our experience with products like SQL*Net and Oracle7 show that network service daemons and information servers benefit from different internal architectures. Most important, since the network itself is the largest bottle-neck, the network listener must be lean and fast, and not carry any extra overhead to its main function.

Today, several hundred web server implementations are in use on the Internet. Ranging from a few lines of Perl code to full-blown commercial products on just about every imaginable operating system, this abundance of available HTTP servers demonstrates two things: first, that it is relatively easy to build a HTTP server, and second, that they will become a commodity very soon. Oracle expects HTTP to become an integral part of all new operating system releases, much like standard FTP or TELNET daemons.

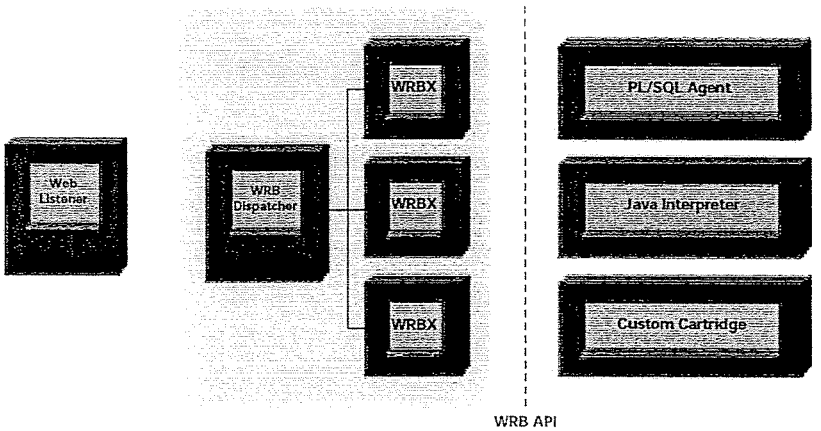
Oracle's mission is to provide an operating system independent development and runtime environment for business applications. As a direct consequence of this objective, we intend to provide support for an increasing number of HTTP engines. Initially, we will provide WRB adapters for the very popular Apache and NCSA servers, but also for other servers such as the Netscape Commerce Server and Open Market's WebServer. Note that these adapters are not provided in the initial release of WebServer 2.0, which only supports the Oracle Web Listener.

Integrating with any HTTP server



The Web Request Broker

Web clients connect to Oracle WebServer 2.0 using HTTP just like any other web server. A request is delivered in the form of a URL, which conceptually is an identifier for a web object (normally an HTML page). The HTTP engine does not make any attempt at interpreting the URL, instead it is immediately handed off to the Web Request Broker for further processing.



The Web Request Broker

The WRB Dispatcher (WRBD)

The WRB dispatcher must decide what type of object is being requested. To do this, it examines the WRB configuration file, which maps virtual directories to WRB Services. For instance, the Web Request Broker may be configured to dispatch all URL requests beginning with `/java` to the integrated Java Interpreter. However, if no suitable WRB Service is found, the request is passed back to the Web Listener and standard processing continues. The WRB Dispatcher performs dynamic load balancing between multiple instances of a WRB Service, and the webmaster can configure each WRB Service to have a minimum and maximum number of instances. An instance corresponds to a process in a UNIX environment.

All inter-process communication is handled by the transport independent WRB protocol. WebServer 2.0 supports standard IPC mechanisms, and future releases will also support other mechanisms, such as for instance OMX, Oracle's CORBA 2.0 compliant Object Request Broker. The advantage of this is clear: WRB Services may be implemented as industry standard distributed objects, enabling true distribution and scalability across multiple nodes.

The WRB Execution Engine (WRBX)

Each WRB Service consists of a common WRB Execution Engine (WRBX) and a shared library, which is dynamically loaded at runtime. By providing an open API to WRBX, Oracle is encouraging partners and customers to integrate their own extensions. Integration is done at a much higher level than CGI or various HTTP server APIs.

The WRB API

The WRB API allows a developer to register three basic callback functions with WRBX: an initialization function, a request handler, and a shutdown function. The WRB API also contains numerous utility functions which may be invoked directly from the three basic callbacks.

The WRB API is being developed in cooperation with Oracle partners and customers, and we are using their feedback to finalize the specification and provide maximum openness and inter-operability between Oracle WebServer and third-party applications.

Building a Cartridge

This section will take you through the steps needed to build your own custom cartridge, or extension, to Oracle WebServer 2.0. Keep in mind that this application development platform is language independent—so you can use any programming language that your operating system supports, as long as your linker can build a shared library from the compiled code. This is a very powerful capability, which allows you to integrate existing legacy applications written in for instance COBOL, as well as develop new cartridges using C or C++. The steps involved in creating a new cartridge are as follows:

1. Create three functions: one to initialize the cartridge, one to shutdown, and one to handle a request.
2. Create an entry point function which takes one parameter of type pointer to a `WRBCallbacks` structure. Fill in the WRBCallbacks structure with your three functions.
3. Link your cartridge with the WRB library to produce a shared library (dynamic link library).
4. Register your library and entry point function in the WRB configuration file using the WebServer Manager.
5. Define one or more virtual directory mappings for your cartridge.

The full specification of the WRB API is available as part of the distribution. Please review it carefully before you attempt to build your first cartridge.

Oracle WebServer 2.0 ships with a sample cartridge called "Hello World!". On UNIX system, you can find the source code in the `ows2/sample/wrbsdk` directory. The file `hello.c` defines the three mandatory functions as follows:

```
WRBReturnCode test_init( WRBCtx, clientcxp )
void  *WRBCtx;
void  **clientcxp;
{
        return (WRB_DONE);
}

WRBReturnCode test_exec( WRBCtx, clientcxp )
void   *WRBCtx;
void  **clientcxp;
{
        WRBClientWrite(WRBCtx, "Content-type: text/html\n\nHello World\n", 40);
        return (WRB_DONE);
}

WRBReturnCode test_shut( WRBCtx, clientcxp )
void   *WRBCtx;
void  **clientcxp;
{
                return (WRB_DONE);
}
```

The first function, **test_init()**, will be our initialization function, which is called once when the cartridge is started. The third function, **test_shut()**, is the shutdown function, which is called once just before the cartridge is stopped. The final function, **test_exec()**, is the guts of the cartridge. It is invoked each time an incoming request is dispatched to this cartridge. As you have probably already guessed, this is a very simple cartridge: nothing is done in the initialization or shutdown functions, and the exec function simply prints a MIME header for an HTML document followed by the words "Hello World".

The next step is to register these three functions as callback function with the WRB Execution Engine (WRBX), which is the program that actually dynamically loads your cartridge at runtime. This is done in the function **testentry()**:

```
WRBReturnCode testentry (WRBCalls)
WRBCallbacks *WRBCalls;
{
        WRBCalls->init_WRBCallback = test_init;
        WRBCalls->exec_WRBCallback = test_exec;
        WRBCalls->shut_WRBCallback = test_shut;

        return (WRB_DONE);
}
```

Finally, you must compile and link your program to produce a shared library which can be loaded by the WRBX at runtime to instantiate your cartridge. The following shows how the shared library is built on Sun Solaris 2.4, but it will differ on just about every operating system:

```
$ cc -c -o helloworld.o -g -I${ORACLE_HOME}/ows2/wrbsdk/inc helloworld.c

$ cc -g -xs -L${ORACLE_HOME}/ows2/wrbsdk/lib -o helloworld.so -G helloworld.o
  -lnsl -lm -lsocket -ldl -laio
```

Please use the Makefile included in the **hello** directory as a template for your operating system. Once the shared library has been successfully built, the next step is to register the cartridge with the Web Request Broker. You can do this with the WebServer Manager HTML user interface. The configuration file in the distribution already has the Hello World cartridge registered as follows (note that **/ora/7.3.2/** will be replaced by whatever your **ORACLE_HOME** is):

```
[Apps]
;
; APP   Object Path                                  Entry Point  Min   Max
; ===   ===========                                  ===========  ===   ===
OWA     /ora/7.3.2/ows2/lib/libndwoa.so              ndwoadinit   0     100
SSI     /ora/7.3.2/ows2/lib/ndwussi.so               ndwussinit   0     100
JAVA    /ora/7.3.2/ows2/lib/libjava.so               ojsdinit     0     100
HELLO   /ora/7.3.2/ows2/sample/wrbsdk/ helloworld.so testentry     0     100

[AppDirs]
;
; Virtual Path          APP     Physical Path
; ============          ===     =============
/ssi                    SSI     /ora/7.3.2/ows2/sample/ssi
/hr/owa                 OWA     /ora/7.3.2/ows2/bin
/ir/owa                 OWA     /ora/7.3.2/ows2/bin
/owa_dba/owa            OWA     /ora/7.3.2/ows2/bin
/java                   JAVA    /ora/7.3.2/ows2/java
/sample/wrbsdk/hello    HELLO   /ora/7.3.2/ows2
/sample/ssi             SSI     /ora/7.3.2/ows2/sample/ssi
/sample/java/run        JAVA    /ora/7.3.2/ows2/sample/java

[AppProtection]
/owa_dba/owa            Basic(Admin Server)

[SSI]
EnableLiveHTML   = TRUE
ParseHTMLExtn    = FALSE
EnableExecTag    = TRUE
ExtensionList    = html shtml lhtml
[JAVA]

CLASSPATH        = /ora/7.3.2/ows2/java/classes.zip:/ora/7.3.2/ows2/java/oracle.zip
                   JavaCacheTimeout = 86400
```

Each line in the [Apps] section is used to define one cartridge. The columns are:

- Cartridge name. A unique name assigned to the cartridge.
- Shared library name. The full path to your application.
- Entry-point. The name of the function that the WRBX will call to register your callbacks.
- The last two columns contain the minimum and maximum number of cartridge instances (threads or processes) that the Web Request Broker will start.

Each line in the `[AppDirs]` section defines a virtual directory path and a working directory for the cartridge. The columns are:

- Virtual directory. When the Web Request Broker sees a HTTP request for an object in this directory path, the request will be dispatched asynchronously to the corresponding cartridge (next column). Note that multiple paths may be defined for the same cartridge.
- Cartridge name. This must correspond to one of the cartridges defined in the `[Apps]` section.
- Working directory for the cartridge.

Each cartridge may optionally have its own section containing parameters, which are accessible through the WRB API. For instance, the Java Interpreter uses the `CLASSPATH` variable to determine where to look for Java classes.

That's all there is to it. This is the way Oracle provides cartridges to extend the base functionality of Oracle WebServer 2.0, and this is the way that 3rd parties and customers can do the same. You will notice that the Java Interpreter, PL/SQL Agent, and LiveHTML are implemented this way. Thus nothing prevents you from integrating any information management system, be it relational databases or legacy systems running on mainframes. If it can be accessed through any programming language on the WebServer node, it can be fully integrated through a custom cartridge and run as a WRB Service.

WebServer SDK

Oracle WebServer 2.0 includes two types of cartridges: system cartridges that perform a well defined function, and programmable cartridges that act as interpreters or runtime environments for applications. An example of a system cartridge is the Verifone VPOS (Virtual Point Of Sale) cartridge, which allows web applications to perform electronic payment transactions using a vast number of payment protocols over the Internet (Note: the VPOS cartridge is not included with Oracle WebServer 2.0, but you can get more information about it at `http://www.verifone.com`). Examples of programmable cartridges are the PL/SQL Agent, Java Interpreter, and LiveHTML, which are all covered in the following section.

The PL/SQL Agent

The PL/SQL Agent enables web applications to be developed using Oracle7 stored procedures written in PL/SQL. This is a tremendous advantage if you are building a data-driven application and your data happens to be stored in an Oracle7 database. Interacting with database objects is much easier in PL/SQL than any other language. It also has the benefit of being very scalable and completely portable between all operating systems supported by Oracle.

From a developer's perspective, the PL/SQL Agent is no different than the popular "Web Agent", which shipped with Oracle WebServer 1.0. Both enable developers to build data-driven, dynamic HTML pages with personalized content using Oracle7 stored procedures. Due to the new architecture in Oracle WebServer 2.0 with the Web Request Broker, the PL/SQL Agent is roughly an order of magnitude faster than the previous version. This is mainly due to the fact that each instance of the PL/SQL Agent stays connected to Oracle7 between requests, and thus does not need to establish a new database connection each time a stored procedure needs to be executed. Note that the Web Agent implemented through CGI is still bundled with WebServer 2.0 for backwards compatibility purposes, but its use is not recommended since the new architecture is so much faster.

The PL/SQL Agent performs two major tasks which enable a stored procedure to generate a dynamic HTML page. The first is to launch the procedure and pass in any parameters that were part of the HTTP request. The translation of a HTTP request to PL/SQL is not obvious, and we will therefore explore this in greater detail in a moment. The second task performed by the PL/SQL Agent is to provide a method for any stored procedure to actually deliver "output". In addition to performing these tasks, the PL/SQL Agent includes the PL/SQL Web Toolkit: a collection of packaged procedures, functions, and utilities to enhance productivity and ensure valid HTML output.

Launching a Procedure

For a request to be handed off to the PL/SQL Agent, the leading directory path must be found in the WRB configuration file. When you install WebServer 2.0, the following directory mappings are pre-defined:

```
/hr/owa                 CWA     /ora/7.3.2/ows2/bin
/tr/owa                 CWA     /ora/7.3.2/ows2/bin
/owa_dba/owa            CWA     /ora/7.3.2/ows2/bin
```

So any request starting with for instance the directory path `/tr/owa` will be recognized and handed off to the PL/SQL agent. Lets look at a sample URL:

```
http://www.oracle.com/tr/owa/tr.splash
```

This URL consists of three different parts: the first part is simply a pointer to the webserver (`http://www.oracle.com`). The second part is the directory path (`/tr/owa`) which has an entry in the WRB configuration file. At this point we know that we will be handing off the request to the PL/SQL Agent. What follows after the directory path (`/tr.splash`) is interpreted as the name of a

stored procedure (the leading slash is ignored), so in this example the PL/SQL Agent will attempt to run a procedure named `tr.splash`. The actual procedure is named `splash` and is part of the public interface to a package named tr, hence the dot-notation.

Connecting to the Database

The Oracle7 server has a higher level of security than flat files managed by an operating system. Users are authenticated using a wide range of mechanisms ranging from simple passwords to fingerprint scanners, smart-cards, and distributed protocols such as DCE and Kerberos. The PL/SQL Agent also needs to authenticate itself to gain access to a distinct schema of objects and storage space within the database. It also needs to have access to all the details required to actually establish a permanent connection: typically a set of environment variables (`ORACLE_HOME` and `ORACLE_SID` for UNIX systems) and perhaps a SQL*Net connect descriptor. All of this information is referred to as a DCD (Database Connection Descriptor), and all of the DCDs known to the PL/SQL Agent are stored in a configuration file named `owa.cfg` (DCDs where known as Web Agent Services in WebServer 1.0). Care should be taken to prevent unauthorized access to this file, since it contains sufficient information to gain access to the database. DCDs are defined in the PL/SQL Agent's configuration file, but they are activated by creating a directory mapping in the Web Request Broker's configuration file.

At the lowest level, Oracle's internal API allows us to establish a connection to an Oracle7 Server using SQL*Net in one step, and actually logging on to the database in a separate step. Since each HTTP request needs a new session, we are technically logging on and off between requests, but this is still extremely fast since the Oracle7 Server connection is maintained between requests.

The token immediately preceding the keyword owa in the directory path is interpreted as the name of a DCD. For instance, `/tr/owa` would cause the PL/SQL Agent to look for a DCD named tr. If such a DCD does not exist, an attempt is made to use the default DCD. If no procedure is found or if the default DCD has not been created, the request will fail.

In Oracle WebServer 2.0 one or more instances of the PL/SQL Agent will be running for each active DCD. Incoming requests will share the same DCD if they are using the same directory path, and authentication of the user will be handled at the HTTP level by the Oracle Web Listener. You will typically configure the Oracle WebServer to have at least one public DCD and several protected DCDs requiring user authentication. When a user has authenticated with a username and password, the username is available to the stored procedure being run, but it is up to the application developer to implement security.

Passing Parameters

The HTTP protocol supports several methods of passing information from the client to the server along with a request. The POST method passes parameters in a stream (standard input) directly to the Web Request Broker (or a CGI program). The GET method tags parameters onto the end of a request. For instance, if two parameters where to be passed to the **tr.splash** procedure, the complete URL would look like this:

```
http://www.oracle.com/tr/owa/tr.splash?name1=value1&name2=value2.
```

The parameters are encoded and passed as name/value pairs with & as a delimiter. If you were developing a web application from scratch, you would have to parse this "query string" and decode the parameters yourself. Fortunately, the Web Request Broker does this automatically for all cartridges, and the PL/SQL Agent converts the parameters to PL/SQL procedure parameters. Furthermore, you don't have to care about what method is used to pass the parameters, they are simply delivered as parameters to your procedure. To complete this example, the URL above with the two parameters would be translated into a PL/SQL procedure call like this:

```
begin
    tr.splash ( name1 => 'value1', name2 => 'value2' );
end;
```

This parameter passing notation may not be familiar to you even if you are an experienced PL/SQL developer, but naming each parameter this way essentially enables us to not worry about the order they appear in. For instance, this procedure could actually have been defined like this:

```
create package tr as
    procedure splash ( name2 in varchar2 := null, name1 in varchar2 := null);
end;
```

Note that all parameters have default values (null). This ensures that the procedure will be found even if one or more parameters are missing. One additional feature is worth mentioning: if you have a large or variable number of parameters, the PL/SQL Agent allows you to use PL/SQL tables of varchar2 as input parameters. Consider an electronic storefront. The user will first search through the product catalog "putting" items into an electronic shopping cart. When the order is ready to be processed, an HTML purchase form with a variable number of line-items needs to be submitted from the browser. If the input fields for each line-item have identical names in the HTML form (e.g. re-occurring **item_partno** and **item_quantity**), the following procedure could handle the request:

```
create package shop as
    type vc2arr is table of varchar2(255) index by binary_integer;
    procedure buy ( item_partno in vc2arr, item_quantity in vc2arr);
end;
```

We have now covered the basics of how an Oracle7 stored procedure is launched by the PL/SQL Agent and how input parameters are passed from the client browser to the procedure. The next section explains how to generate the resulting dynamic HTML using the PL/SQL Web Toolkit.

Generating HTML with the PL/SQL Web Toolkit

The toolkit consists of a number of packages that are installed into each schema corresponding to a DCD when the DCD first is created. The two packages named **HTP** and **HTF** (HyperText Procedures and HyperText Functions) make up the core of the PL/SQL Web Toolkit and provide a low-level one-to-one mapping between HTML 2.0 and PL/SQL procedures and function. At the lowest level, the procedure **htp.print** plays a central role, in that it is ultimately used either directly by the developer or indirectly by a higher-level procedure or function to put something into the result, which is transmitted back to the client. Although the result will usually be a HTML document, any MIME type is supported, including user-defined types. Here is an example of a simple procedure:

```
create procedure hello is
begin
    htp.print('Hello World');
end;
```

The **htp.print** procedure simply puts whatever parameter you pass to it on a page. A shortcut procedure is available with the name **htp.p**. You can of course add HTML tags manually, or you can use the functions and procedures in the PL/SQL web toolkit that encapsulate HTML for you. The following two procedures generate the exact same result, **hello1** simply prints the result exactly as desired, whereas **hello2** uses the toolkit procedures and functions:

```
create procedure hello1 is
begin
    htp.p('<title>Hello Sample</title>');
    htp.p('<h1>Hello Sample</h1>');
    htp.p('Hello World, this is <b>easy</b>');
    htp.p('<hr>');
    htp.p('<a href="http://www.oracle.com">Visit Oracle on the Web!</a>');
end;

create procedure hello2 is
begin
    htp.htitle('Hello Sample');
    htp.p('Hello World, this is ' || htf.bold('easy'));
    htp.line;
    htp.anchor('http://www.oracle.com','Visit Oracle on the Web!');
end;
```

The second example demonstrates a powerful way of bundling multiple HTML tags into widgets. The **htp.htitle** procedure generates both a title and a header tag. This concept can easily be expanded to higher levels, so you could

for instance define a common procedure to set your document title and header, background bitmap, text and link colors, and navigation toolbar. In addition to increasing productivity, this technique automatically brings a consistent look and feel to your applications.

Java Interpreter

For a general introduction to Java, please visit `http://java.sun.com`. This section assumes general knowledge of Java as a technology, although a very brief introduction follows. Java is an objective-oriented language strongly resembling C++ but with a few restrictions that strongly improve application robustness, the most important being no explicit pointer manipulation and no multiple class-inheritance.

Java allows you to develop three types of classes: normal Java classes which are used as building blocks or abstract classes, runnable classes, known as Java applets, which can be downloaded by a Java-enabled web browser, and Java applications, which are standalone programs. Oracle WebServer 2.0 allows you to integrate the last type, standalone applications, into your website using the Oracle Java Interpreter. Oracle WebServer 2.0 (and 1.0 for that matter) can also act as a repository for Java applets to be downloaded to clients, but that is a different topic not covered here.

The distinguishing attribute of a Java application compared to an applet or other class, is that it contains a public method (function) called `main`. This method acts as the entry-point to the application. Oracle WebServer 2.0 ships with a few sample Java applications, of which the simplest looks like this:

```
class HelloWorld {
    public static void main (String args[ ]) {
            System.out.println("Content-type: text/html\n\n");
            System.out.println("Hello World!");
            System.out.println("Hello World!");
    }
}
```

Anything printed to the standard output stream from a Java application is simply delivered back to the client.

Extending Java

Oracle WebServer also contains the Oracle Java Web Toolkit, three Java packages to extend the capabilities of the standard Java interpreter. These are `oracle.html.*`, which contains objects for dynamic HTML generation, `oracle.rdbms.*`, which contains objects for database access, and `oracle.plsql.*`, which contains objects for PL/SQL access. These packages must be explicitly imported by any application that intends to use them.

Executing PL/SQL from Java

The Java Interpreter interfaces to the Oracle7 Server by running PL/SQL packages or standalone PL/SQL procedures and functions. Each package or procedure must have a Java package wrapper, which is a Java class containing methods to call that package's procedures and functions. Standalone procedures and functions are all wrapped in a single wrapper. Once you have identified or created the PL/SQL packages your applications needs, you can create the package wrappers automatically by running the `pl2java` utility as follows:

```
pl2java [flags] username/password[@connect-string] packagename...
```

This utility creates a Java wrapper class for each package or procedure given as an argument to the command. When your application is run, it creates an instance of this class to interface to the package. The following is a summary of steps required to write a Java application which interacts with an Oracle7 Server through PL/SQL:

1. Create an object of type Session to handle the database connection. All of the operations performed during this session must be called from within this object.
2. For each PL/SQL package used in the application, you must create one instance of the packagewrapper subtype created for that PL/SQL package by the `pl2java` utility.
3. For each parameter of a PL/SQL package, you must create an instance of the Java variable that matches that parameter.

The following example uses a packaged function named **`employee.count_employees`** to get a tally of people in a given department. For the sake of simplicity, all error and exception handling has been left out.

```
import oracle.rdbms.*;
import oracle.plsql.*;

public class EmployeeReport {
     public static void main (String args[]) {
               .
               int tally = get_count("SALES");
               .
     }

     int get_count( String deptName ) {

     Session.setProperty("ORACLE_HOME", "/user/oracle");
     Session.setProperty("TNS_ADMIN", "/user/oracle/network/admin");

     Session session;
     session = new Session("scott", "tiger", "HR_DB");

     // Create a new instance of Employee package
     Employee employee = new Employee(session);

     // Create object to encapsulate PL/SQL parameter
     PDouble       pEmployeeCount;
```


ORCL000776

```
    // ADD: call Employee package to count the number of
    // employees in the department
    pEmployeeCount = employee.count_employees(pDeptName);
    return (int)pEmployeeCount.doubleValue();
  }
```

Generating HTML from Java

A Java application can generate a dynamic HTML page. Essentially, anything written to standard output by a Java application run by the Java Interpreter will be returned to the client browser. The "Hello World" example described earlier demonstrates this. Similar to the PL/SQL Web Toolkit, the Oracle Java Web Toolkit allows you to generate HTML without having to hardcode the actual HTML tags.

To generate dynamic HTML from within Java, you create various objects that use the interface **`IHtmlItem`**. All classes that generate dynamic HTML implement this interface, which has two simple methods: **`toHTML`** and **`print`**. Both of these methods produce the content of the object as HTML, but **`toHTML`** returns it as a string, whereas **`print`** sends it to system output. Therefore, you use **`toHTML`** when you want the resulting HTML to be further processed by another method and use **`print`** when you want to output it. In effect, you build your Web page in a buffer with **`toHTML`** and flush the buffer with **`print`**.

The **`oracle.html`** package provides a standard set of classes based on HTML2, HTML3, and popular browser-specific extensions. You are not limited to these, however. You can easily create your own customizable HTML classes by deriving them from the **`CompoundItem`** or **`Container`** classes. The **`oracle.html`** package also has the intelligence to generate output that is optimized for the browser at hand. For example, a browser that does not support tables will get table data in the form of pre-formatted strings. Here is a simple example:

```
import oracle.html.*;

public class HelloWorld {
    public static void main (String args[]) {

            // Create an HtmlHead Object titled "Hello World!"
            HtmlHead hd = new HtmlHead("Hello World!");

            // Create an HtmlBody Object
            HtmlBody bd = new HtmlBody();

            // Create an HtmlPage Object
            HtmlPage hp = new HtmlPage(hd, bd);

            // Adds a simple string "Hello World" in this page
            bd.addItem("Hello World!");

            // Print out the content of this Page
            hp.print();
    }
}
```

LiveHTML

One of the most popular features of the Web is known as *Server Side Includes* (SSI). Certain documents can be parsed by a web server before sending them back to the client, and they may contain "include" directives instructing the web server to insert information into the document on the fly. Examples of use are page-hit counters, timestamps, etc. Server Side Includes literally bring otherwise static HTML pages to life, hence the name LiveHTML. Oracle WebServer 2.0 supports SSI+, a common extension to NCSA's original specification.

LiveHTML code is formatted as SGML comments, so that it is ignored should the file ever find its way to the browser unparsed. The format for LiveHTML codes, therefore, is the following:

```
<!--#command tag1="value1" tag2="value2" -->
```

The tags are arguments to the commands, most of which actually only accept one of the possible tags. The possible commands and their associated tags are as follows:

`config`. This command sets parameters for how the file or script is to be parsed and therefore is normally the first LiveHTML command in a file. The possible tags are:

- `errmsg`. This specifies the error message that is sent to the client if an error occurs while parsing this document. Here is an example:

  ```
  <!--#config errmsg="A parse error occurred in the Hockey_Pass file"-->
  ```

- `timefmt`. This specifies a date format. LiveHTML files frequently include timestamps. The conventions follow the strftime library call supported in most versions of Unix, even if the WebServer is not running on a Unix platform.

- `sizefmt`. This specifies the format used when displaying a file size. The possibilities are bytes, which gives the absolute size in bytes, and abbrev, which gives the size in kilobytes or megabytes as appropriate.

- `cmdecho`. This specifies whether non-CGI scripts subsequently executed have their output incorporated into this HTML page. The possible values are ON and OFF. ON specifies that the output is included. The default is OFF.

- `cmdprefix`. Specifies a string that will be prepended to each line of the script output.

- `cmdpostfix`. Specifies a string that will be appended to each line of the script output.

`include`. This command specifies that a file is to be included in the generated HTML page at this point. The file can any of the following:

- another LiveHTML file like the current one.

ORCL000778

- a regular HTML file.
- an ASCII file. which of these it is determined, as usual, by the extension.

The possible tags are:

- `virtual`. This gives a virtual path to the file. The directory mappings for virtual paths are set by the WebServer administrator using WebServer Manager.
- `file`. This gives a pathname relative to the current directory. References to parent directories or uses of absolute pathnames are forbidden.

`echo`. This gives the value of an environment variable. This variable is either one of the standard CGI environment variables or one of the LiveHTML extensions, which currently are all standard Server Side Include variables. There is only one tag, var, and it must be present. It provides the name of the variable. The LiveHTML environment variables are as follows:

- `DOCUMENT_NAME`. The current filename.
- `DOCUMENT_URL`. The virtual path to this file.
- `QUERY_STRING_UNESCAPED`. If the client sent a query string, this is an unescaped version of it, with all shell-special characters escaped with \. DATE_LOCAL The current date and local time zone in the format specified by the most recent config timefmt command.
- `DATE_GMT`. Same as the above, but in Greenwich mean time.
- `LAST_MODIFIED`. The last modification date of the file, given in the format specified in the last config timefmt command.
- `fsize`. This produces the size of the file in the format specified in the most recent config filesize command. Tags are the same as for include.
- `flastmod`. This produces the last modification date of the file in the format specified in the most recent config timefmt command. Tags are the same as for include.
- `exec`. This is the command to execute a script. The tags specify whether or not the script is CGI.
- `cmd`. This specifies a non-CGI script. Execution is passed to the Operating System, and the given string is parsed as though it were entered at a command-line interface. The full path of the script must be given. The non-CGI environment variables specified under echo above can be referenced. Whether the output of the script is included in the HTML page that the parser outputs, is determined by the most recently executed config cmdecho.
- `cgi`. This specifies a CGI script. The value given will be the virtual path of the script. URL locations are automatically converted into HTML anchors.

Summary

First generation webservers have fundamental shortcomings in the areas of scalability, management, and extensibility. Oracle is leveraging its core competency in these areas to deliver a next-generation product capable of addressing these issues and making extremely large websites with hundreds of thousands of users and terabytes of information a reality. Oracle WebServer 2.0 lays the foundation and provides a first glimpse of new emerging technology, which will be fully developed in forthcoming releases. After all, this is Internet Technology, and we are finding our product life-cycles being reduced from years to months.

Oracle WebServer 3.0 is scheduled for release in late 1996, and we expect the Web Request Broker API to be fully developed by that time. Oracle is actively working with partners and developers to provide a solid platform for application development. Other plans for WebServer 3.0 include making extensive use of Oracle7 for website configuration management, including user authentication, optionally storing content with automatic replication to multiple sites, more cartridges, inter-cartridge processing, and full support for all major HTTP listeners. It is going to be an exciting year—be sure to keep an eye on our public website, `http://www.oracle.com`, for upcoming news, releases, and events.

Magnus Lonnroth
Senior Product Manager
Oracle Internet Server products

Terminology and Acronyms

Term	Description
Apache	An increasingly popular HTTP server on the Internet
API	Application Programming Interface
CGI	Common Gateway Interface
daemon	Detached process waiting for incoming service-requests over the network
FTP	File Transfer Protocol
HTML	HyperText Markup Language
HTTP	HyperText Transfer Protocol
Java	An object-oriented language for the Web
JavaScript	A language developed by Netscape for use in the Netscape Navigator and the Netscape Commerce Server through NSAPI
JDBC	ODBC for Java
LiveHTML	Currently same as Server Side Includes.
MIME	Multi-purpose Internet Mail Extensions
NCSA	National Center for Super-computing Applications
NSAPI	Netscape Server API
PL/SQL	Oracle's procedural language for stored procedures
SNMP	Simple Network Management Protocol
SQL	Structured Query Language
SQL*Net	Oracle's network protocol which runs on top of most networks
SSI	Server Side Includes
SSL	Secure Sockets Layer
TELNET	Terminal Emulation Program
URL	Uniform Resource Location
VRML	Virtual Reality Markup Language
W3C	World Wide Web Consortium
WRB	Web Request Broker
WRB API	An API for developing custom cartridges
WRBD	Web Request Broker Dispatcher
WRBX	Web Request Broker Execution Engine

ORCL000781

Exhibit 49

ORACLE® White Paper

Oracle and The Internet

Using Oracle in the Emerging Electronic Infrastructure of Business and Society

July 1995

Part A30972

Oracle and The Internet

Part A30972

July 1995

Author: Jack Haverty

© 1995 Oracle Corporation

All rights reserved.

Printed in the US.

This document is provided for informational purposes only and the information herein is subject to change without notice. Please report any errors herein to Oracle Corporation. Oracle Corporation does not provide any warranties covering and specifically disclaims any liability in connection with this document.

Oracle and SQL*Net are registered trademarks, and Oracle7, PL/SQL, Oracle Internet Server, Secure Network Services, Oracle Media Server, and Oracle Text Server are trademarks of Oracle Corporation.

All other company and product names mentioned are used for identification purposes only, and may be trademarks of their respective owners.

Oracle and The Internet

Using Oracle in the Emerging Electronic Infrastructure of Business and Society

July 1995

ORCL00937781

Oracle and The Internet

Using Oracle in the Emerging Electronic Infrastructure of Business and Society

The Internet is creating the new infrastructure for business and society. Oracle complements the Internet, communicating information processing with industrial strength mechanisms.

Introduction

The Internet is arguably one of the hottest high-tech topics today. It has emerged from the academic and research world to become a de facto component of the "Information Highway." Corporations and organizations have begun to use the Internet to understand both the capabilities and limitations of the new technology, and how it can be best applied in their own operations.

Many of the limitations of the current Internet technology parallel the obstacles encountered in the IS environment over the last decade, which drove the creation of Oracle information management technology and its acceptance as the premier choice for production quality information systems.

This White Paper describes Oracle's perspective on the synergy between the Internet and Oracle technologies, and our vision for the Internet as an element of the emerging new infrastructure for business and society.

The Internet is a new environment. Everyone is exploring how to best use it for their needs.

The Challenge of The Internet

There is a frenzy of activity within the Internet, as people from all walks of life and corporations from all industries rush to get involved. Participation in this modern "Gold Rush" is being viewed as a crucial step for all sorts of organizations who want to establish their position in the new electronic world.

Savvy organizations are participating with a recognition of the state of the technology and future directions. Like pioneers in a new land, organizations in cyberspace don't really know what is over the next mountain range, but they do know that there will be surprises along the journey. They can see other pioneers, and learn from observing mistakes.

ORCL00937783

They also learn from successes - their own and others, and will form relationships and practices that may be used for decades to come.

Exploring the Territory

Early pioneers on the Internet replicate old physical mechanisms in the new electronic world. Innovators explore the unique capability of networks to make their Internet presence stand out from the crowd.

In the exploration of the Internet, most organizations focus on two goals to figure out how to apply this new technology in their business machinery: getting up the learning curve, and identifying how the Internet affects their own particular situation or business. Pioneers in a new land have always sought a better life, and similarly business and individuals on the Internet and the broader "cyberspace" seek ways to improve their existence.

On the Internet, that situation has typically spawned two kinds of exploration. First, electronic mechanisms are implemented that simply mirror current physical mechanisms - electronic malls, newstands, and so on. Second, new ideas are nurtured and tried, exploiting the differences between the old and new environments and the innovations made possible by new technology.

For example, many organizations have established online presences to disseminate their "message", basically adding the Internet as another means for distributing standard marketing materials. Others have recognized the different cost structure and electronic "reach" of cyberspace, where business techniques become viable that would otherwise be infeasible - such as a "consignment shop" for tidbits of information each of which costs a few dollars.

The World Wide Web is a simple but powerful implementation of client/server systems using a highly open, distributed, cooperative approach.

The technology which has triggered all of these activities is the World Wide Web - a set of techniques, protocols, technology, and software, which enables users in a wide range of business and society to actually use the Internet, both as consumers and also as producers of information and services that are accessible worldwide.

In the new environment of cyberspace, innovation and high-quality content is what allows an organization to "stand out from the crowd". In the physical world, "location, location, location" is the mantra of business; in the virtual world, where every location is a mouse-click away, content has become the key to success.

World Wide Web Basics

In the Web technology, a basic client/server architecture underlies all activities. Information is stored on Web Servers. Web clients, such as Mosaic software running on a PC, Mac, or Unix workstation, request information by specifying the "name" of the information in a request transmitted across the network to a server. The server simply transfers the requested information back to the client - in much the same way that

ORCL00937784

the various network file systems and servers operate on many networks today. The client formats the information and displays it on the user's screen.

Information is typically organized into "pages", each of which is up to several screenfuls of material. Pages can contain "hyperlinks" that allow a user to readily "jump" to another page, which may actually be stored on another server many miles away.

The information in pages is typically stored as text files, containing the information encoded in HTML - HyperText Markup Language. With HTML, pages can be readily created with formatted textual information freely mixed with multimedia content including graphics, audio, and even video, as well as the "links" that enable a user to "browse" from page to page to page across the Internet..

In addition to pages which are stored in files on web servers, the web technology provides a mechanism for dynamic information to be included in pages. The Web server in such cases uses a standard interface, called the Common Gateway Interface (CGI), to execute a separate program which obtains the dynamic information and returns it formatted as HTML.

The *Common Gateway Interface* is a key standard. It enables the power of database mechanisms to be easily applied on the Internet.

Using this interface, it is becoming common practice for pages to include information from many kinds of sources. The external program might trigger a camera to capture an image and pass it back to the Web Server for delivery to the Client. On the Web there are pages which deliver images of highway interchanges, building lobbies, or simply interesting views from somewhere else in the world.

With the CGI interface, it is becoming common to extract information from databases. The program executed through the interface is simply a client of the database, and performs a query or stored procedure operation to obtain the appropriate information and format it as HTML for return to the user.

These core activities of the Web are illustrated in the diagram on the next page, where Oracle7™ is being used to supply information to a Web server for delivery to some client across the Internet.

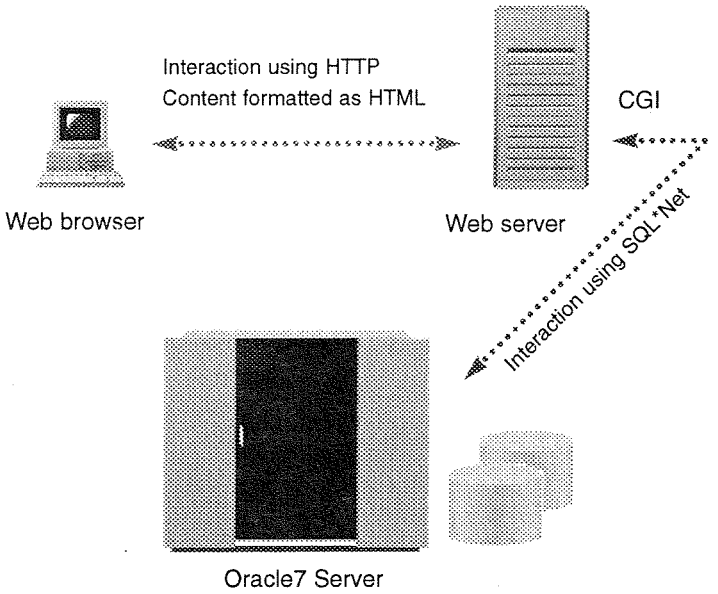
Evolution of the Web

Despite the Web's youth, the content of information on it has already started to evolve, as the pioneers try various ideas and learn from experience.

The genesis of the Web was as a mechanism for sharing static information - scientific reports, proposals, papers, and the like, which were pre-

ORCL00937785

The Web began as a simple mechanism for storing information. It has already evolved into a platform for serious distributed applications.



pared 'offline' and loaded onto a Web server. In effect, the Web was simply a distributed "file system", which made it easy to publish material and cross-reference it, and equally easy to "browse" and view the information. The Web provides a way to share static information, i.e., material prepared offline and placed onto the Web by a largely manual process.

The second phase of Web practice was enabled by the inclusion of the CGI mechanism mentioned above, which enables pages to be dynamic and interactive. Some or all of their content is created on-the-fly, in response to input from the user. If you browse the Web today, you'll probably observe that 5 or 10% of the pages you visit are dynamic and/or interactive. Information pages that display current weather conditions at a location specified by the user, or price quotes for specified stocks, are examples of rudimentary dynamic and interactive information.

The content on the Web has now evolved to a third phase, in which collections of cooperative static and dynamic information and interactive pages create services and applications such as electronic storefronts. If you browse the Web today, you'll find maybe about 1% or less of the pages you visit implement some kind of application such as an electronic store. For people familiar with databases, client/server systems, and business applications, this will look like familiar ground, because it is.

ORCL00937786

The Role of Oracle Technology

Driving the Evolution

The need for applications, in business and society, is driving web technology today. Database technology was driven by similar needs over the last decade, and is readily applicable to the web.

Behind the evolution of the Web, there are fundamental forces driving innovation. The progression from static to dynamic information and interactivity has occurred for the same reasons that SQL and other database technology appeared in the IS environment - people need ways to specify how to use information in order to apply computer and networking power to performing tasks in business, education, and everyday living.

Merchants want to sell their products and interact with their customers and partners via the network. Governments want to provide services to their citizens. Publishers want to deliver information online in a way that their customers can find what they're seeking. They all need the same kinds of facilities in the Internet and Web that have been in use for years in the more private environments of corporate and organizational networks and computing facilities.

The three phases described previously reflect these needs - not just for a mechanism to access information, but for an infrastructure for using information. Services and applications are the target. Early Web implementations provide access to content, both for producing and consuming that content, just as network file systems have made content accessible to a distributed community. Applications which use, manipulate, input, modify, analyse, and apply that content to real-world tasks are the next step.

Many of the limitations and gaps in today's web technology are obstacles to large-scale use as a real infrastructure.

Databases and applications constructed "on top" of them have created Information Systems and business mechanisms using the underpinnings of computers, file servers, networks, and the like in today's corporate world. Databases and associated applications are readily applicable to the Web and Internet to make the same kind of jump from accessing data to using it as information and services.

Overcoming Obstacles

Organizations in all areas of business and society have been learning by implementing live systems on the Internet and exploring the power and potential of the Internet and World Wide Web for use in their various activities. At the same time, they have observed the limitations and immaturity of the technology and products. As is true of most technologies emerging from the "pipeline" of R&D, the Internet and World Wide Web have several inherent weaknesses. For example, they lack mechanisms which the research community deferred to the vendors for productization of the technology, such as billing, accounting, administra-

ORCL00937787

tion, management and other core business and operations facilities.

Scale is another issue - systems which can handle a relatively small and technically astute community are not necessarily applicable to a massive population of everyday users. Approaches that are viable in early pilots and trials will need to be implemented for production with technology and products that are viable in large-scale, core operations in order to become infrastructures of electronic commerce and society. Fortunately, a lot of that proven technology is applicable today, and extensions of current technology to support the new environment of the Internet and Web is straightforward and natural.

Oracle relational database technology provides several powerful techniques for implementing web systems with proven products and familiar techniques.

The Roles of Oracle7 In The Web

Using relational database technology within the Web is straightforward because of the open nature of the Web technology and its architecture, the inherent nature of the Web as a primitive form of distributed database, and the maturity of relational database technology with its own set of open industry standards. As mentioned above, the standard Common Gateway Interface provided on a Web Server can be readily used to interact with Oracle7. There are a variety of situations in which this kind of architecture is used:

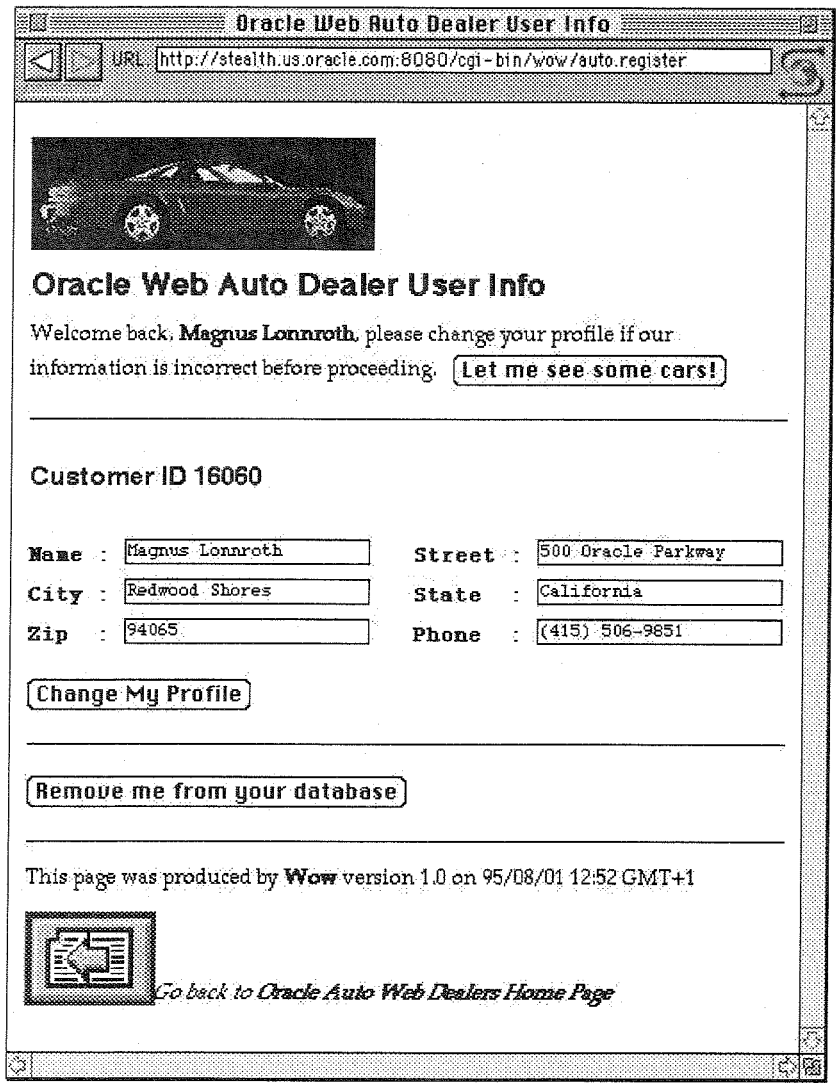
- Data which is managed in Oracle7 databases is made available to information consumers via the Web, using the power of the relational database to provide dynamic Web pages, in which the facilities of Oracle7 are used to create information on demand from raw data and application logic.
- Applications which are used on the Web use the programmability of Oracle7 to implement functions of the application itself. Inventory status, airline or hotel reservation availability, interactive pages for requesting product literature, for registration to attend an event or purchase a product. Much more elaborate applications are feasible as well. Such "transactions" can also be readily integrated with existing systems for order entry, billing, and other back-office mechanisms, since they are likely already functioning in a database environment.
- Information, even "static" information such as Web pages prepared offline, can be stored in Oracle7, and delivered in response to client requests from the Web. By managing such information with a database, all of the traditional Information Systems facilities of a database are applicable -

ORCL00937788

such as the ability to track the ownership, revision history, and changes to a page, and the ability to manage and synchronize content by transaction techniques and distributed database mechanisms.

To illustrate the synergy of these capabilities, the Oracle World Wide Web Interface Kit, which was released on the Internet in February1995 includes an example application, an "Auto Dealership" electronic commerce mockup. This demonstration application uses a design in which every page is purposely dynamic, created "on-the-fly" by stored proce-

Entire electronic storefront applications can be implemented as single PL/SQL packages stored in the Oracle7 Server.



A screen shot from the Auto dealer demo released as part of the Oracle World Wide Web Interface Kit

ORCL00937789

dures. There are no files storing HTML as in early Web servers, but HTML is created on demand in response to user requests arriving via the Internet.

Leveraging the Power of Networked Databases

In all of these situations, advanced techniques of today's Oracle7 products can be applied to address the requirements for large-scale, high-performance business systems. For example, multiple Web servers could be interfaced to powerful parallel-processing systems running Oracle in order to handle large numbers of transactions.

Oracle's mature products for distributed systems can be readily used to create production-quality, scalable, innovative web presences.

Similarly, multiple Oracle7 servers could be configured and operate as a single logical database, performing remote access to retrieve information from wherever it is stored, or using replication to maintain multiple copies of information in many places. Web servers and Oracle7 interact through the Web server's Common Gateway Interface and the standard Oracle Call Interface.

Because of the Oracle networking architecture, Web servers can be located separately from the database; several Web servers could, for example, be distributed across the Internet, each acting as a front-end to data and applications running on a corporate centralized Oracle7 server. Using SQL*Net and Secure Network Services™, these Web servers and database servers could be interconnected using the Internet, with assurance of security.

With Oracle and related technology and products from thousands of partner vendors, our customers can create a single virtual Web presence with dynamic, interactive content that stands out from the crowd because of its innovation and quality. They can manage that content with powerful and well-proven tools and techniques. They can deploy the content in a wide variety of configurations, from single-site to multiple distributed physical servers. They can evolve from pilot to production to large-scale operation in a natural and proven way, leveraging from the techniques developed in the relational database systems of the last several years as they have undergone similar growth, and leveraging the database skills and expertise that are usually already present in many organizations.

Product Directions

Oracle makes products for managing information. Oracle's customers use these products in large-scale, "mission critical" business environments, as the core of their electronic infrastructure.

The product suite, built around the core of "server" components, has

ORCL00937790

As the Internet and World Wide Web evolve the new open information infrastructure, Oracle products provide world-class production implementations.

been built to address and solve all of the kinds of real-world issues discussed above. Oracle-based information systems are industrial-strength. They scale to accommodate thousands of simultaneous users, with many gigabytes, terabytes, and probably before long petabytes of data, and many MIPS of computing power dispersed across collections of physical machines. Data is secure, manageable, and safe. Business applications are reliable, fast, and efficient. That has been our business focus for almost twenty years, and Oracle is recognized today as the leader in that industry.

Oracle World Wide Web Interface Kit

Oracle's initial focus in Internet and World Wide Web technology has been the application of current products to this electronic environment. The Oracle World Wide Web Interface Kit has provided guidance and assistance for our customers to couple existing production Oracle database technology with the Internet and World Wide Web, and begin creating pioneering content and applications.

The Oracle World Wide Web Interface Kit provided the industry's first guidelines on integration of database and web technology, and on construction of web applications using RDBMS tools and technology.

The Interface Kit included an initial version of a PL/SQL-based package which enabled use of Oracle7 as a datastore and platform for Web-based applications. That package, called "WOW" to reflect Web-Oracle-Web integration, has been used by pioneering Oracle customers on the Web to create the kinds of dynamic, leading-edge applications mentioned earlier.

Oracle Internet Server

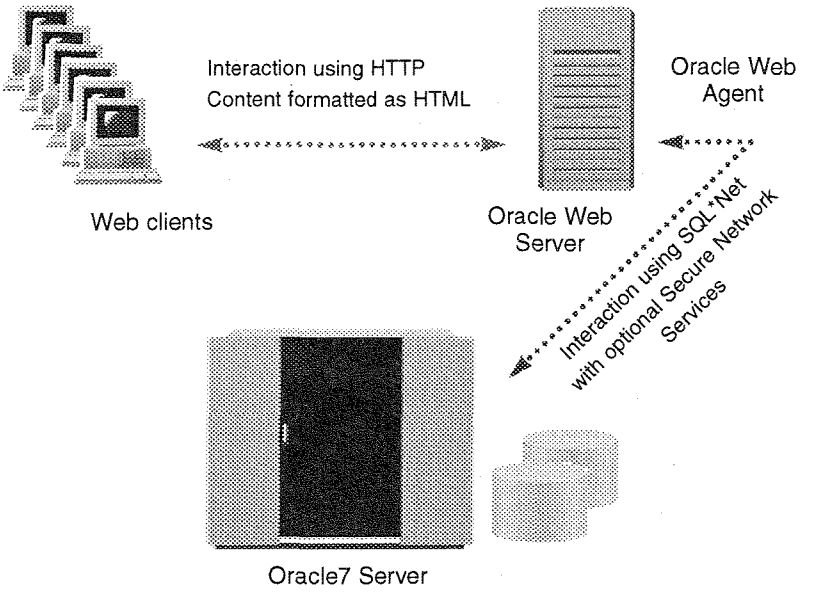
The Oracle Internet Server™ is a standard Oracle product, available on many platforms and supported worldwide. It contains three components:

- A full Oracle Web Server, conformant to current standards and comparable to other vendors' products.
- The Oracle Web Agent, which is a descendant of the early "WOW" gateway in the Interface Kit and couples the Web Server to Oracle7 using CGI mechanisms. The Web Agent provides an object-oriented, user-extensible framework for producing dynamic HTML pages using Oracle7 stored procedures written in PL/SQL.
- The Oracle7 Server - the RDBMS itself, licensed to be used as a standalone embedded component of the Oracle Internet Server, or interfaced with standard distributed database techniques to other Oracle7 installations as a part of an organization's overall IS environment.

ORCL00937791

These three components are integrated in the Oracle Internet Server product with installation, configuration, and operations as a composite system. Each component can be used independently as well. A Web Agent and Oracle7 can be coupled to an existing Web Server, or a Web Server and Web Agent can be coupled to an existing Oracle7 installation.

The Oracle Internet Server provides an integrated product solution for online applications servicing thousands of users with high amounts of information.



A customer can use the Oracle Internet Server as a packaged, fully integrated, complete platform for establishing a web presence, and readily "scale up" when needed by adding any of the elements of the Oracle7 technology, including products from Oracle as well as thousands of other vendors.

As with any Oracle7 client, the Web Server and Web Agent will be able to operate within the same computer as Oracle7, or they can be deployed separately, interconnected by any kind of network using Oracle's SQL*Net, Secure Network Services, and other networking products.

The Oracle Internet Server is available worldwide as a member of Oracle's product suite, with all of Oracle's product and services including WorldWide Support, Education, and Consulting. Oracle customers can choose this product for their Internet activities with the assurance that it will be supported on many different platforms, backed up by expertise in essentially every country. In addition, Consulting services

ORCL00937792

are available to assist with any phase of a project from initial study through design, implementation and ongoing maintenance. The support of the thousands of partners in Oracle's Business Alliance Program with their products and services that synergize with Oracle's components complete the picture.

Oracle Internet Server Directions

Integration of web server and RDBMS components provides the high performance and load capacity needed for large-scale, production online applications.

The Common Gateway Interface is a convenient way to interface a Web server with an RDBMS. However, it has limitations, especially with performance and scaleability. Oracle's product efforts in this area are focused on providing tighter integration of Web server and database components, removing the bottleneck that the convenient approach includes.

With the next release of the Oracle Internet Server, the Oracle Web server and Oracle7 components will be interfaced directly in the Web Agent, rather than through the CGI mechanisms. This tighter coupling will provide the same functionality for HTML-based content and applications as with the CGI-based approach, and upward compatibility for content created earlier. However, due to the tighter coupling, performance will be improved, especially scaleability.

With this enhanced implementation, the Web Server will be a native "client" of Oracle7, able to sustain the higher levels of performance and user-load needed for production-class Web applications.

CGI-based configurations will continue to be supported, so that other Web Servers and other CGI-based external programs can be used, but these of course will not benefit from the tighter integration. For pilot and early production systems, such CGI-based implementations are quite usable. However, they can be readily migrated to the more integrated configuration when the operational needs of performance and scale dictate.

In addition to the tight coupling to achieve production-quality performance and scale, Web Agent and Web Server mechanisms will take advantage of RDBMS approaches for management of the information they contain. Much of the technology associated with databases can be readily applied to Web systems. For example, today's approach for storing material in the Web is to use files, which works well in small-scale situations but becomes unwieldy as the amount of information increases. Storing and managing such material in a database is a primary focus of Oracle's product activities. Integration of other Oracle technology, such as the text management facilities of Oracle ConText™, Oracle Text Server™ and Oracle Media Server™ components, will make those

ORCL00937793

technologies also readily applicable to Web applications.

Some (R)evolutionary Ideas

With the synergy of these technologies and the backing of Oracle products, a wide range of innovative and powerful applications can be envisioned and built with real-world production deployment as the target. As an example, let's take a quick look in two concrete areas of "electronic commerce", namely sales and advertising, to illustrate how innovative content can be created using Web and database techniques.

Selling Electronically

Database and information handling technology offers the opportunity for the innovative approaches to online mechanisms that apply the technology to business needs.

On the Internet today, the first phase electronic commerce for the sales process is already visible. Electronic shopping malls, storefronts, and even "industrial parks" have already been built and are operating, using whatever techniques are available to achieve needed levels of reliability, security, and functionality. Such systems reflect the first techniques used by pioneers in cyberspace - simply taking a familiar physical mechanism and creating its counterpart in the electronic world.

New technology creates new opportunities, and new capabilities that didn't exist in the physical world. Consider for example a sales cycle in which a consumer wants to buy a product - such as a cordless telephone. That consumer might visit the electronic mall in the same way he or she would visit a physical mall, and such activities are common today on the Web. But in the world of cyberspace, other techniques are viable, and hold the potential for revolutionizing commerce.

Instead of an electronic mall, the customer might visit a "personal shopping" service - another storefront which maintains a database by searching the Web and researching product information, sales terms, stock availability, and other aspects of interest to the consumer. The consumer might find in that shopping service a set of reviews of cordless telephones by independent test labs, product literature from various vendors, discussion logs from online forums of owners of various products, and much other related information. Armed with such information, he or she might select a product, and even select a vendor to purchase it from - or might in fact use the information gathered to visit a local physical mall and purchase the product. Cyberspace does not replace the real world, but rather complements it.

This "re-engineering" of commerce might continue. The shopping service might be focused on home electronics, and might give the consumer a free disk with a "Music Organizer" application - a simple application for cataloging and indexing the user's music collection in a data-

ORCL00937794

base residing on the user's home computer. When that user connects again to the shopping service, the service's database and user's database might interact, to form a profile of that user's interests and create personalized pages of equipment, concerts, compact discs, and other material for that particular user.

All of this can be done today, with today's technology and today's products.

Advertising Electronically

Continuing this example, perhaps the shopping service announces to the user that a new music store has opened, specializing in music from all genres and artists, and offering highly personalized service. When the user connects to that store's electronic presence, he discovers that the store is heavily subsidized by sponsors, with each page containing one or more small "ads" - small logos or banners scattered on various pages - a situation becoming common today on the Web.

However, after browsing through the store and seeing a number of pages, the user notices an unusual situation. All of the ads are actually interesting - every artist mentioned is a favorite, every concert promoted is at a nearby site. In fact, all of the advertising information, and even the product content of the pages, is personalized to that specific user. The advertising might include a reminder from the user's insurance company, or dentist, or auto mechanic, or ... let your imagination fill in the blanks!

Under the covers, there are databases at work, and innovative applications. This kind of application can be created today, using today's technology and today's products.

Databases are obvious components of a server, but they are also highly applicable elsewhere, such as at the client.

Client Equals Server

Such innovative applications are based on advancing to the next step in integration of database and networking technologies. Integrating a Web server with a database server is a fairly obvious step. But consider what happens when a database is integrated with a Web client....

Today, many clients have extremely primitive databases embedded in the client itself. The popular Mosaic browser for example has a "history" file and a "hot list" where a user can jot down interesting places on the Web in order to find them again. Such facilities are barely recognizable as a "database", but in fact that's what they are.

Consider the possibilities if a real database were present — in or accessible from the client computer on the Web. My "browser" might now be able to remember not only where I've been, but what I've looked at. It

ORCL00937795

might provide me with a "search" facility that can handle queries like "find me the place I visited a few weeks ago where cordless telephones were being discussed." With Oracle7, TextServer, ConText, and other such technologies applied on both the "client" and "server" sides of the Web, all sorts of innovative possibilities occur.

This next step, integrating database servers with client machines in addition to servers, effectively transforms the client into a server, and permits innovative applications such as the sales and advertising programs described earlier.

We've had client/server systems for years. Oracle pioneered Client-Agent-Server approaches in 1994 with Oracle Mobile Agents. Client Equals Server is a natural next step - and it's available today with today's products and creative content implementations.

What Has Oracle Done?

All of the frenzy around The Internet, and all of the issues outlined above, apply to Oracle as well as our customers, partners, and competitors. Oracle has a unique position in this new frontier, with a role not only as a vendor but also as a user of the new technologies and media of cyberspace.

Oracle as a User

Oracle has a unique background for the new world of networks and information in the Internet.

Oracle's internal worldwide network is one of the largest of any software company.

Oracle has been heavily involved in the "networking" industry, not only with the extensive product set focused around SQL*Net, but also with Oracle's own internal corporate network. The Oracle Internet is a clone of "The Internet." It uses the same kinds of routers, LANs, circuits, WAN services, and other such components as are used in the public Internet, and has been built out by Oracle's own networking staff worldwide and used since 1989 to provide internal communications for Oracle's operations.

Oracle has been actively involved in the recent leading edge as well. On the foundation of our internal network, Oracle staff, from all parts of our organization, have built our own internal private Web. The Oracle Web now contains servers and their information and service "content" located at Oracle offices in many countries, built and operated by local in-house personnel, and by Oracle staff around the world.

The components of the Oracle Internet Server have been used internally in these efforts. Hands-on, operational experience is not only helpful in providing internal services - it is mandatory for obtaining the experience with the technology necessary to develop viable infrastructure-quality products.

ORCL00937796

From that basis of internal experience, in 1994 Oracle brought its own corporate Web presence online on the public Internet, not only to provide corporate and product information to the user of the net, but also to actually deliver products over the new electronic channel. From its starting point in the US, that Web presence has been extended to provide service "closer" to customers beginning with Europe as the next step. Oracle is a major "cyberbusiness", using the network to interact with our customers in all aspects of our products and services activities.

Oracle has been a hands-on user of Internet technology for many years, and has applied that experience to product development.

Oracle has used the Internet for years, and has provided products for using databases in client/server and distributed configurations over all kinds of networks, including the Internet. We are "webbed" internally and externally. We are users. We see the same opportunities as our customers see. We see the obstacles and issues. We feel the same pain as we learn lessons in the trenches of new technology. We apply the learning to products for the Internet environment.

Oracle as a Vendor

As the major vendor of open database technology, Oracle is familiar with all of the issues of the World Wide Web.

As a vendor, much of Oracle's core product technology is already applicable in Internet implementations as discussed earlier. One way to view the World Wide Web at this stage of its evolution is as an extremely primitive implementation of a distributed database system - and wildly successful despite its limitations. Applying Oracle technology injects world-class database capabilities into the Web, and enables the creation of innovative and impressive applications for use over the network by a worldwide user community.

In the industry and standards organizations, thousands of technologists at companies worldwide are working to advance the core technologies of the Internet and the World Wide Web, filling the gaps and extending the technology based on the experiences from the operations in the current Web. Oracle is supporting and working closely on the key activities in this area that are forming the next generation of standards for the Web and Internet.

Much of the technical standardization comes from the Internet Engineering Task Force, the technical arm of the Internet Society, of which Oracle is a Founding Member. Technical developments specific to the Web are focusing in the World Wide Web Consortium (W3C), a joint and global effort of the Massachusetts Institute of Technology and INRIA in France. Application of the technology to business systems, and in particular to the creation of true Electronic Commerce, is a focus of groups such as CommerceNet, an industry consortium which Oracle sponsors and supports, driving the deployment of electronic commerce on the Internet.

By involvement in these key activities that lead from research and technology development through standardization and deployment in business systems, Oracle's efforts couple with the experience we gain from being users, and drive directly into our product development efforts.

Oracle Internet Server an Applications Platform for the Web

The Oracle Internet Server provides database information processing capabilities as a platform suitable for all kinds of applications.

The target of all of these efforts is to rapidly advance the state-of-the-art of Internet technology and real-world systems such as Electronic Commerce. By taking two well-developed technological bases - the Internet and Oracle7 - we create a platform in which all sorts of innovative applications can be readily and easily constructed and deployed, changed and evolved.

This "Platform" operates as an open system in the environment of the Internet and Web where "open" is more than a buzzword - it is mandatory for survival. By providing solutions for several key issues that today's Internet and Web users are debating, we expect that databases in general, and the Oracle platform in particular, will become the platform of choice for implementation of Web-based applications - not just for a simple Web experiment, but rather as the core of the emerging electronic commerce systems worldwide.

Industrial Strength

Today's Web servers are immature both as technologies and as products, and unproven for use in 7x24 large-scale production environments. Issues of robustness, load capacity, and service levels to users are key factors for a professional business system.

The early trials and pilots of web pioneers need industrial strength systems in order to go in production.

Database servers are available on a wide range of platforms, from personal machines to massively parallel machines. With an integrated Web server acting as a front-end and interface to a database, accessing it as a normal client, industrial strength implementations are feasible.

Since the Web server is just a client of Oracle7, it can reside on the same machine as Oracle7. But it can also reside on another machine, accessed across a network by using SQL*Net in a client/server configuration. The network might in fact be the Internet, but it could also be an organization's private network - even one that doesn't use the Internet's TCP technology, such as a NetWare or LU6.2 corporate enterprise net. Organizations might even have dozens of Web servers scattered across the Internet, all acting as front-ends to a centralized Oracle7 system where the data and applications are actually housed.

Within the Internet, the responsiveness seen by a user is strongly related to the "distance" between the user and the server. Adding bandwidth to alleviate performance problems is a popular panacea, but unfortu-

ORCL00937798

nately has real economic impacts, and content with high-quality multimedia is likely to grow faster than the underlying networks, or budgets, can handle.

Placing the server closer to the user means faster response times, but requires many servers to be used at many locations, and therefor introduces problems of management and synchronization of all of the copies of the data. An organization desiring a global presence in the Web might, for example, operate servers in many locations around the world, mirroring a "master" server. With Web servers alone, that creates a nightmare of maintaining complex distributed files. If the content is dynamic, such as a reservation system, such distributed schemes are probably impossible with pure Web technology. But such issues were recognized years ago in the database industry, and solved with products available today. With Oracle7, replication technology and distributed transaction mechanisms provide an easy solution, simply by applying database products in the web and Internet.

In other words, Industrial Strength is not as much of an issue as it might look. By applying and integrating even today's Oracle technology with today's Web technology, the path to large-scale, reliable, robust, manageable Web systems becomes much more clear.

Information Infrastructure

One of the most common complaints heard about the web is the difficulty of finding things; an obvious use of database approaches.

Imagine taking all of your corporate data, printing it onto single sheets of paper, and then tossing that huge stack of paper up into the air in a large room. Now find something that you know is there. Now suppose that the large room is actually your building's lobby, and the 9 AM rush is in progress, with paper being scattered and mixed by thousands of feet. Now browse the Web, and see if it feels any different.

In the world of Oracle customers, these kinds of situations are avoided by using database technology and all of its ancillary tools, procedures, services, and mechanisms. With Web information actually stored in a database, and with applications that have ready access to that database, an information infrastructure for your Web information can pop into existence overnight. Existing tools can be used, current staff skills are applicable, and the Web is an adjunct to existing information systems.

If you want to use the Web to collect customer information, just write an application that creates an appropriate fill-in form and puts the data gathered into tables. If your customers are having trouble finding specific information within your organization, write a Web application that provides a searching interface to your data. If you'd like to track where visitors go within your Web material, just enhance your application to

ORCL00937799

log whatever information you need into appropriate database tables.

Security

Security in the Internet is a major issue, and is getting major attention to make the Internet technology ready for real-world uses like electronic commerce. The bulk of these efforts in defining technology and standards address the linkages between Web clients and servers. Standards are being enhanced to include encryption and authentication, procedures are being deployed for identifying users, and techniques for performing paper-based activities, such as signing documents, are being developed in the electronic world.

All of these efforts are proceding rapidly, and will emerge as products from many vendors as the pioneer organizations and individuals, vendors and consumers, of the Internet become ready to settle down and do "real" electronic commerce. Products from many vendors, including Oracle, will support and embrace standards that enable "real" use of the net.

Security concerns inhibit current use of the web for real world tasks. Similar concerns drove the development of database security mechanisms, which are ready to apply in web implementations.

Such efforts are useful and necessary, but they address only a part of the problem of security. Control of access to data for example, namely who is permitted to use which information, are still very primitive capabilities, not even in the realm of the standards activities which largely focus on interactions "on the wire" between clients and servers, as opposed to actions within applications or servers themselves.

Such issues are the same kinds of issues that the database industry and users have encountered in the evolution of database technology, and many of the solutions that exist now are applicable through the Web as well. User roles, table views, and other capabilities in Oracle7, for example, provide powerful techniques for controlling access to information. Facilities such as triggers are very useful in providing mechanisms such as audit trails. Business rules can enforce the structure of data. Trusted Oracle7 can be used to provide a high level of assurance of security. Distributed database techniques can be used with Secure Network Services to achieve security and assurance that the data is intact and private even if it happens to traverse the Internet's worst neighborhoods.

Migration and Coexistence

Organizations have an enormous investment in their current IS environment, and in existing expertise and installed production systems. The techniques discussed above provide an obvious path where existing systems can be used and extended to the Internet and Web. An organization's data can be accessible through the Web, without losing its acces-

ORCL00937800

sibility to existing users, and without requiring tedious and expensive processes to export, import, and synchronize multiple versions of the same information.

Through current IS tools such as forms-based applications, and through leading-edge technologies such as Oracle's Mobile Agents for wireless access, the Web can be a front-end to a database, and the data can be accessed and manipulated from traditional as well as Web-based interfaces.

The web is new, and most organizations struggle with how to integrate it with their current business systems - a task greatly simplified by the Oracle Internet Server.

The Web also provides techniques for creating powerful hybrid systems. For example, a Web server and database application might deliver a page which contains some information - such as a product description - and a button for the user to order a product. That button could send a request to the server which results in an Oracle Forms document being returned to the client. The browser, from any of the many vendors of browsers today, could receive that document, recognize it as a Forms document, and pass it to the local runtime Oracle Forms application.

In this scenario, the Forms document might be an organization's existing "order entry" application, providing the Web user with the ability to order a product without requiring a separate application to be developed for the Web. Or it could be an application written for the Web, taking advantage of the powerful features of the Forms environment to make a Web presence which "stands out from the crowd" when compared with traditional Web applications pages.

Competing Frontiers

The web is huge and growing, but it is only one of several new electronic media, all of which any organization may need to utilize.

As successful and prominent as the Internet and Web are today, they are not alone as leading-edge technologies. Other technologies, such as Interactive TV, CD-ROM, and even Virtual Reality, provide similarly new worlds in which business and societal tasks can be easily envisioned. Pioneering in one of these areas is a struggle for many organizations; pioneering in all of them may be fatal for all but the largest organizations.

The techniques discussed above, in particular for delivering content and services as applications operating in the server environment, provide a powerful approach for mitigating the risk of choosing the wrong frontier, and minimizing the expense of pioneering in several. Oracle's strengths in classical relational databases, in multimedia video and interactive television, in document and text management, and in evolving technologies such as object-oriented systems, are creating a "multi-frontier" platform on which Oracle's customers can participate in one or more of these frontiers.

ORCL00937801

The electronic storefront described above might provide one "door" for customers on the Web, and another "door" for customers using their home televisions and set-top boxes. Both would access the same data, the same information, the same application.

The future "Information Superhighway" is likely not to be the Internet. Nor will it be Interactive TV. Nor a proprietary online service. Nor any other single approach. It will be some combination of all of these, and maybe some we haven't seen yet. Pioneers didn't pick one frontier and build civilization - they picked all of them and stayed where life worked best.

With that vision, hybrid systems become attractive as well. A Web application, for example, might provide an electronic storefront for an auto dealer accessible to anyone's home computer - but when the customer wants to "test drive" a car's 10-minute video clip, it could be delivered over an ISDN channel to that customer's home entertainment center.

The frenzy of exploration results from organizations seeking both the efficiencies of electronic mechanisms and the innovation which new technology makes possible. Oracle technology and products have a natural role in providing the industrial-strength implementations for taking new electronic mechanisms to production use. In addition, the marriage of networking and relational database technologies provides a fertile ground for innovation, applying the techniques and capabilities of database systems to create innovative business approaches that exploit the technological advances inherent in The Internet.

Looking Forward

Oracle's product efforts for the Internet are focused in two areas. The first is simply to enable and encourage the use of existing Oracle products, applying them to the Internet and Web, to apply the database capabilities that are sorely needed in Web information and applications.

The second focus is the extension of the existing products, and creation of new ones, that implement the new technology with the same traditional "industrial strength" nature of Oracle's core products, working with the community to define and create the technology and standards to fill in the gaps in the current approaches of the Internet.

With this focus, Oracle customers can get involved in their own exploratory efforts in the Internet. Using RDBMS facilities, people with database experience can pursue innovative approaches, working from familiar and stable ground. Networking people can similarly explore database technology, working from their own territory. These pioneer-

ORCL00937802

ing efforts can proceed with the comfort of having a well-developed suite of technology and products that exist today, as well as those that will follow, and provide scalability and "industrial strength", along with support and ancillary products and services available from Oracle and its many partners worldwide.

The Internet and the World Wide Web mark the emergence of the network industry into the information realm. This event makes it now possible to apply all kinds of Oracle technology using the Internet as a foundation.

Much Oracle technology can be readily applied to the Internet. Searching services are of course obvious - Oracle7, TextServer, and related components have a natural "fit" into the Internet. Other systems, such as Oracle's MediaServer and video-on-demand technology, can be applied - if the Internet can handle the load! Hybrid systems are possible - where the Internet is one of several delivery mechanisms that can be used to connect customers with information and services.

Our goal is simple:

> *Oracle will be the customers' preferred choice of platform for information storage and processing within the emerging electronic infrastructure of commerce and society worldwide.*

Watch for lightning bolts as we and our customers bring the two industries and technologies of database and The Internet together.

ORCL00937803

ORACLE®

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
USA

Worldwide Inquiries:
415.506.7000
Fax 415.506.7200
http://www.oracle.com/

Printed on recycled paper with soy-based inks

To offer our customers the most complete and effective information management solutions, Oracle Corporation offers its products, along with support, education, and consulting, in more than 90 countries.

Oracle and SQL*Net are registered trademarks, and Oracle7, Oracle7 Workgroup Server, Oracle Internet Server, and PL/SQL are trademarks of Oracle Corporation.

All other company and product names mentioned are used for identification purposes only, and may be trademarks of their respective owners.

Copyright © Oracle Corporation 1995
All Rights Reserved
Printed in the USA

Part #: A30972

Exhibit 50

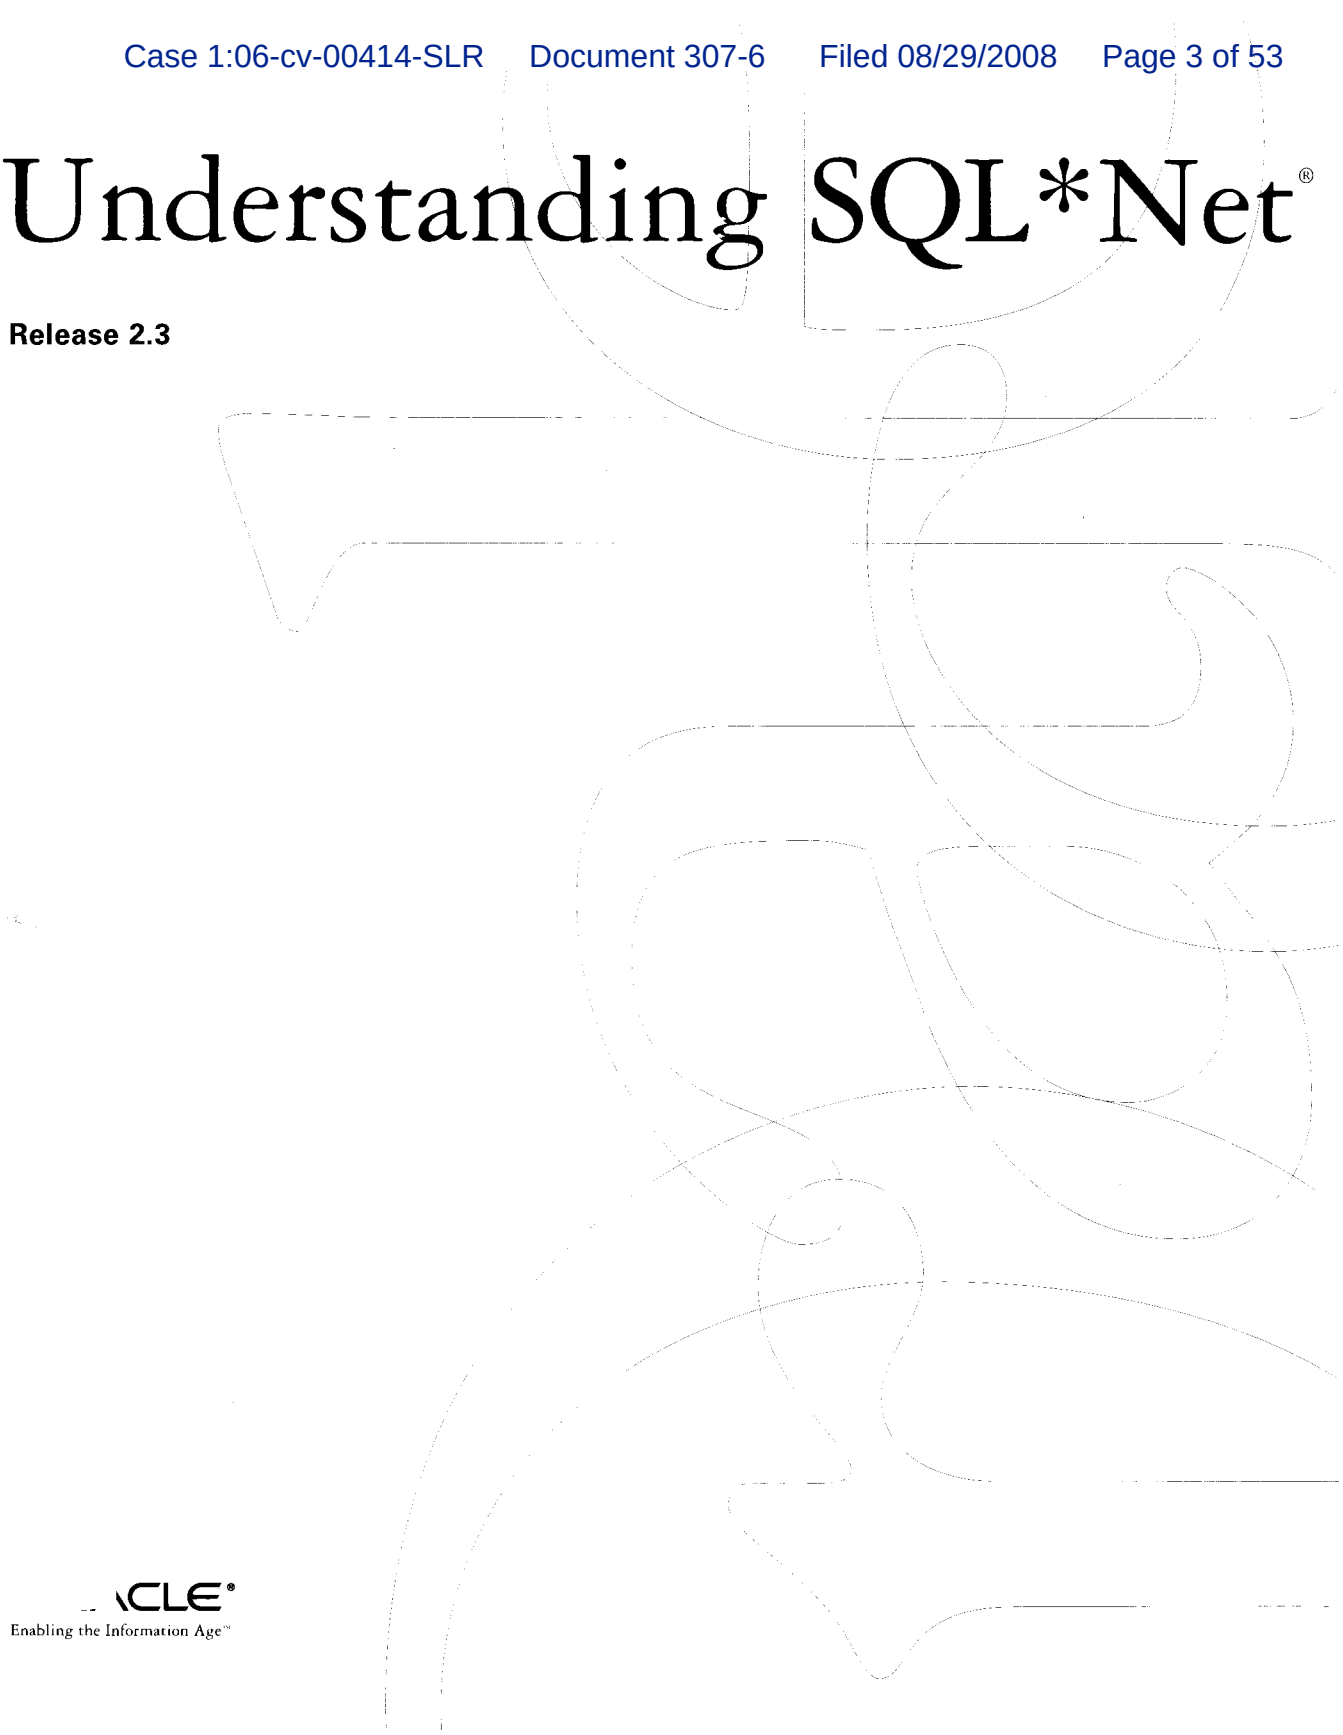
Understanding SQL*Net®

Release 2.3 (A42484-1)

ORACLE®

Understanding SQL*Net®

Release 2.3



ORACLE®
Enabling the Information Age™

Understanding SQL*Net®

Release 2.3
Part No. A42484–1

ORACLE®

Understanding SQL*Net, 2.3

Part No. A42484–1

Copyright © 1992, 1993, 1994, 1995, 1996 Oracle Corporation

All rights reserved. Printed in the U.S.A.

Primary Authors: Larry Neumann, Sandy Venning, Laura Ferrer
Contributing Author: Shelly Dimmick Willard
Contributors: Romeo Baldeviso, George Chang, Craig Clow, Peter Chu, Carl Dellar, John Graham, Jack Haverty, Mark Hill, Mark Jarvis, Mike Kanaley, Ed Miner, Joe Pistritto, Jay Rossiter, Cyril Scott, Deb Smith, Dave Stowell, Steve Viavant, Rick Wessman, David Wong, Lelia Yin

This software was not developed for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It is the customer's responsibility to take all appropriate measures to ensure the safe use of such applications if the programs are used for such purposes.

This software/documentation contains proprietary information of Oracle Corporation; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited.

If this software/documentation is delivered to a U.S. Government Agency of the Department of Defense, then it is delivered with Restricted Rights and the following legend is applicable:

Restricted Rights Legend Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of DFARS 252.227–7013, Rights in Technical Data and Computer Software (October 1988).

Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

If this software/documentation is delivered to a U.S. Government Agency not within the Department of Defense, then it is delivered with "Restricted Rights", as defined in FAR 52.227–14, Rights in Data – General, including Alternate III (June 1987).

The information in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Oracle Corporation does not warrant that this document is error free.

Oracle, SQL*Net, SQL*Forms, SQL*DBA, SQL*Loader, SQL*Menu, and SQL*Plus are registered trademarks of Oracle Corporation.

Oracle MultiProtocol Interchange, Oracle Names, Oracle Network Manager, Secure Network Services, Oracle Server Manager, and Oracle7 are trademarks of Oracle Corporation.

All other products or company names are used for identification purposes only, and may be trademarks of their respective owners.

QA
76.9
.D3
U528
1996

CHAPTER

5 Using SQL*Net

This chapter describes how to use SQL*Net release 2.3 after it has been configured. It includes:

- starting and testing the network configuration
- common errors during testing
- using the Listener Control Utility
- initiating a SQL*Net connection
- distributed database management
  - database links
  - synonyms
  - snapshots
  - copying data between databases
- operating system-authorized logins

Testing the Network Configuration

Note: This section assumes that you are not using Oracle Names 2.0 with the Dynamic Discovery Option enabled. If you are using the Dynamic Discovery Option, please refer to the *Oracle Names Administrator's Guide* for information about starting the network.

Configure the network using the Oracle Network Manager. Once the configuration files are on the destination machines, each component can be started and tested. The preferred sequence for testing the network is to:

- start and test each Oracle Names Server (if included)
- start and test each listener
- start and test each Interchange (if included)
- test each client type

Start a Names Server

The STARTUP command of the NAMESCTL utility loads the Names Server into memory and tells it to begin executing. At startup, the Names Server loads its configuration, loads its data, then becomes available to answer requests.

```
NAMESCTL> STARTUP
```

Test a Names Server

To test that a Names Server is operating correctly, use the PING command. Following are two ways to PING the server LABRADOR in the US.ACME domain. Alternatively, you can use TNSPING (described on page 5–5).

```
NAMESCTL> PING LABRADOR.US.ACME
or
NAMESCTL> SET SERVER LABRADOR.US.ACME
NAMESCTL> PING
```

You can also PING multiple Names Servers by using a single command, for example:

```
NAMESCTL>PING HUEY.UK.ACME DUEY.UK.ACME LOUIE.UK.ACME
```

PING responds with the time it takes to contact the Names Server and return an acknowledgment.

If PING fails, make sure the Names Server has been started. If it has been, double-check the configuration in Network Manager, especially the defined address of the Names Server in question.

ORCL01808684

Start the Listener

From each listener's node, use the Listener Control Utility, LSNRCTL, to start each listener. In command line mode, to start a default listener (which is a listener defined in a LISTENER.ORA file with the alias LISTENER):

```
LSNRCTL START
```

To start a non–default listener:

```
LSNRCTL START listener_name
```

LSNRCTL will display a status message indicating that it has started successfully. After LSNRCTL has started the listener, it has no further control over the listener. Check that all expected SIDs for that listener are listed in the services summary in the status message. For more information on the Listener Control Utility see the section "Using the Listener Control Utility" later in this chapter.

Test the Listener

To test the listener, initiate a connection from a client in the same community as the listener to any active database controlled by that listener.

The simplest test uses SQL*Plus as follows:

```
SQLPLUS user/password@service_name
```

The *service_name* may be found in the TNSNAMES.ORA file, an Oracle Names Server, or a native naming service such as NIS or DCE's CDS. For more information on testing SQL*Net from a client, see the section "Initiating a SQL*Net Connection" later in this chapter.

If there are no clients in the same community as the listener, you must start an Interchange before testing the listener.

Repeat these steps for each listener in the network.

Start theInterchange

Use the Interchange Control Utility, INTCTL, on the Interchange node to start an Interchange. (You cannot run INTCTL from a remote node.) For example, in command linemode the command is:

```
INTCTL START INTERCHANGE
```

or

```
INTCTL START INT
```

INTCTL displays a status message indicating that it has started successfully. For more information on the Interchange Control Utility see Chapter 6, "Controlling the MultiProtocol Interchange," in the *Oracle MultiProtocol Interchange Administrator's Guide*.



Intchg

ORCL01808685

Test the Interchange

To test the Interchange, initiate a connection through the Interchange from a client in one community to a database in another. On the client machine type:

```
SQLPLUS user/password@service_name
```

The *service_name* for the database in the other community may be stored in TNSNAMES.ORA, an Oracle Names Server, or a native naming service such as NIS or DCE's CDS. To ensure that the connection went through the Interchange, type:

```
INTCTL STATUS INTERCHANGE
```

The Connection Manager should indicate that there is one active connection through the Interchange. You can run the STATUS command on the Interchange machine. If the Interchange is listed in the TNSNAMES.ORA file, stored in an Oracle Names Server, or stored in a native naming service, then you can run the STATUS command remotely from any other node.

Test Network Objects Using NAMESCTL

To test that a network object exists, use the QUERY comand. The syntax for this command is as follows:

```
NAMESCTL> QUERY global_object_name type
```

Database service names have the type A.SMD, and database links have the type DL.RDBMS.OMD. The following example shows a query of the database service name BUGSY in the MACS.ACME domain.

```
NAMESCTL> QUERY BUGSY.MACS.ACME A.SMD
```

The QUERY command returns the amount of time the transaction took and information about the network object.

Test Network Objects Using TNSPING

The network administrator can use TNSPING to determine whether or not a service on a SQL*Net network can be successfully reached. The service being contacted can be an Oracle database, an Oracle Names Server, or any other Oracle (TNS) service.

If you can connect successfully to a TNS service using TNSPING, it displays an estimate of the round trip time (in milliseconds) it takes to reach the Oracle service. If it fails, it displays a message describing the error that occurred. This allows you to see the network error that is occurring without the overhead of a database connection.

You invoke TNSPING on the command line as follows:

```
tnsping service_name count
```

On some platforms, such as Microsoft Windows, the interface may be different, but the program accepts the same arguments.

ORCL01808686

- *service name* must exist in TNSNAMES.ORA, Oracle Names, or the name service in use, such as NIS or DCE's CDS. (For service names to be resolved to addresses in a native naming service such as NIS or DCE's CDS, Oracle Native Naming Adapters must be installed on clients and servers. See your Oracle operating system–specific documentation for information on installing and configuring Oracle Native Naming Adapters.)
- *count* (optional) determines how many times the program attempts to reach the server.

If the service name specified is a database name, TNSPING attempts to contact the corresponding network listener. It does not actually determine whether or not the database itself is running. (To do this, use Server Manager to attempt a connection to the database.)

Examples of the Use of TNSPING

Following are some examples of how a network administrator or user can use TNSPING. For example, to connect to a database named SPOTDB, the command:

```
tnsping spotdb
```

produces the following message:

```
TNS Ping Utility for SunOS:  Version 2.2.2.0.0 - Production on
10-Mar-95 10:09:13
Copyright (c) Oracle Corporation 1995.  All rights reserved.
Attempting to contact
(ADDRESS=(COMMUNITY=build_tcp.us.oracle.com)(PROTOCOL=TCP)
(HOST=spot)(PORT=1599))
OK (50msec)
```

To check whether an Oracle Names Server can be reached, use a command similar to the following:

```
tnsping (ADDRESS=(COMMUNITY=build_tcp.us.oracle.com)
(PROTOCOL=TCP)(HOST=fido)(PORT=1600))
```

The address of the Oracle Names Server is located in SQLNET.ORA (in the NAMES.PREFERRED_SERVERS and NAME.PREFERRED_SERVERS parameters) on the client.

> **Note:** Both parameters showing the client's preferred Names Server must be included for backward compatibility with an earlier version of the product.

A message similar to the following will be returned to the user:

```
TNS Ping Utility for SunOS: Version 2.2.2.0.0 - Production on
10-Mar-95 10:09:59
Copyright (c) Oracle Corporation 1995. All rights reserved.
```


ORCL01808687

```
Attempting to contact
(ADDRESS=(COMMUNITY=build_tcp.us.oracle.com)
(PROTOCOL=TCP)(HOST=fido)(PORT=1600))
OK (70 msec)
```

To determine whether the STPRD database can be connected to, and to specify that TNSPING try to connect 10 times and then give up, use the following command:

```
tnsping stprd 10
```

This command produces the following message:

```
TNS Ping Utility for SunOS:  Version 2.2.2.0.0 - Production on
10-Mar-95 10:10:28
Copyright (c) Oracle Corporation 1995. All rights reserved.
Attempting to contact (ADDRESS=(COMMUNITY=build_tcp.us.oracle.com)
(PROTOCOL=TCP)(HOST=spot)(PORT=1599))
OK (290 msec)
OK (100 msec)
OK (70 msec)
OK (70 msec)
OK (60 msec)
OK (70 msec)
OK (70 msec)
OK (80 msec)
OK (180 msec)
OK (340 msec)
```

Below is an example of TNSPING attempting to connect to an invalid service name:

```
tnsping invalid_service_name
```

This attempt produces the following message:

```
TNS Ping Utility for SunOS:  Version 2.2.2.0.0 - Production on
10-Mar-95 10:10:58
Copyright (c) Oracle Corporation 1995.  All rights reserved.
TNS-03505: Failed to resolve name
```

Following is an example of using TNSPING to connect to a name that is valid, but that resolves to an address where no server is located (for example, the server may be down or not started):

```
tnsping testing
```

The following message is returned:

```
TNS Ping Utility for SunOS: Version 2.2.2.0.0 - Production on
10-Mar-95 10:11:42
Copyright (c) Oracle Corporation 1995.  All rights reserved.
```


ORCL01808688

```
Attempting to contact
(ADDRESS=(PROTOCOL=tcp)(HOST=spot)(PORT=1599))
  TNS-12541: TNS:no listener
```

Test All Client Types

Make sure that each client type is tested. It is not enough to test that the Interchange works. If there are several different client types in your network, initiate a connection to a server from each of them. If a connection is unsuccessful, use logging and tracing—the diagnostic tools—to find the cause of the problem. An error stack in the error log may point to the problem. If not, turn on tracing and repeat the operation. You can find information about error logging and tracing in Chapters 1 through 3 of the *Oracle Network Products Troubleshooting Guide*.



Diagnostics

Common Errors During Testing

If you are unsuccessful in bringing up a listener, Interchange, or Names Server, or fail to make a connection to a database, check to see if the cause is one of the following common errors.

- An invalid listener name was typed in the LSNRCTL START command.

  Check your typing. Check the LISTENER.ORA file to ensure that the listener name you are using is valid.

- Files were put in the wrong place.

  Both the listener and the Interchange will indicate that they cannot start because configuration files could not be found. Check your operating system-specific documentation to see that the LISTENER.ORA file has been placed correctly on the listener machine, and the INTCHG.ORA, TNSNAV.ORA, and TNSNET.ORA files are placed correctly on the Interchange machine. (Refer to "Testing the Network Configuration" in this chapter, and to the *Oracle Network Products Troubleshooting Guide* for more information.)



OS Doc



Diagnostics

  If you are using TNSNAMES.ORA files, make sure they were created with Network Manager, and not created by hand. If you are using Oracle Names Servers, make sure they have been started. Also make sure that NAMES.ORA is placed correctly, and that SQLNET.ORA file contains the preferred Names Server parameter. If a native naming service such as NIS is in use, make sure that the appropriate Native Naming adapter has been installed on clients and servers, and that service names have been



OS Doc

ORCL01808689

properly loaded into it. Refer to your operating system–specific documentation for information.

- Address already in use.

  Another process may already be using the address listed in LISTENER.ORA. On some protocols such as TCP/IP, DECnet, and OSI, each network service on a node must use a unique port or socket. On other network protocols such as SPX/IPX or NetBIOS, each network service name must be unique for the entire network. Another network service may be using the same configuration. Contact your network administrator to evaluate whether the network address is available.

- When trying to connect to a database, you may get the message ORA–12203: "TNS:Unable to connect to destination" .

  Use the LSNRCTL utility to start the listener on the server machine. See Appendix C for further information.

- When trying to make a connection from a client, you may get the message ORA–12154:"TNS:Could not resolve service name".

  The service name you requested is not listed in the TNSNAMES.ORA file, or the TNSNAMES.ORA file has been placed incorrectly. See Appendix C for further information.

- When trying to connect to a database, you may get the message ORA–1034: "Oracle Not Available".

  The database is not running on the server machine. A listener alone does not provide a database connection; the database instance must also be started.

- Connect requests that come in too quickly for a listener to handle, and which exceed the listener's backlog (determined by QUEUESIZE parameter in LISTENER.ORA, TNSNET.ORA, and NAMES.ORA), are returned with an ECONNREFUSED error. A client encountering this error returns the message "ORA–12541: No Listener" and the client log or trace files will show the ECONNREFUSED message.

  To correct this problem, follow these steps:

  – Stop the listener.

  – Reconfigure QUEUESIZE in LISTENER.ORA, TNSNET.ORA, or NAMES.ORA to be a larger value (based on anticipated simultaneous connect requests).

  – Restart the listener.

– Try to connect again.

Refer to "Increasing the Queue Size for the Listeners" in Appendix A for information on increasing the number of connection requests the listener can handle.

Note: The default for QUEUESIZE is operating system-specific. Similarly, it is dependent upon the operating system whether the QUEUESIZE parameter will be recognized and be put into effect.

- When attempting to stop the listener, you may get the message TNS–01169: "The listener has not recognized the password." Enter the SET PASSWORD command from within LSNRCTL, and then enter the STOP command to stop the listener.



Diagnostics

Other common errors are listed in Appendix C of this manual. All the error messages generated by SQL*Net, the Interchange, Oracle Names, Oracle Native Naming Adapters, and Oracle Network Manager and their underlying layers can be found in the *Oracle Network Products Troubleshooting Guide*. The *Oracle Network Products Troubleshooting Guide* also contains information about how to interpret log files and how to use the trace facility for troubleshooting purposes.

Using the Listener Control Utility

The network listener establishes listen endpoints on a machine. These listen endpoints are well–known addresses that clients and servers use to connect to a database. The Listener Control Utility, LSNRCTL, is a tool that you run from the operating system prompt to start and control the listener. For SQL*Net use, the general form of the Listener Control Utility is:

```
LSNRCTL command [listener_name]
```

In this syntax:

LSNRCTL	Specifies the name of the tool that controls the listener. In some operating systems, this fixed parameter is case sensitive. If the operating system is case sensitive, enter LSNRCTL in lowercase.
command	Specifies the action to be performed.
listener_name	Specifies the listener name defined in LISTENER.ORA. If *listener_name* is not included, the default value LISTENER is used.

ORCL01808691

You can also issue Listener Control commands at the program prompt. When you enter LSNRCTL on the command line, the program is opened. You can then enter the desired commands from the program prompt. For example, the following command starts the database subagent for a node managed by SNMP support.

```
LSNRCTL> dbsnmp_start
```

Commands for LSNRCTL Information

The Listener Control Utility has three commands that provide information about LSNRCTL itself. These commands are HELP, SET, and SHOW.

HELP

HELP provides a list of all the LSNRCTL commands available. An example follows:

```
iris[331] lsnrctl

LSNRCTL for SunOS: Version 2.3.1.1.0 - Beta on 17-AUG-95 17:15:02

Copyright (c) Oracle Corporation 1994.  All rights reserved.

Welcome to LSNRCTL, type "help" for information.

LSNRCTL> help
The following operations are available
An asterisk (*) denotes a modifier or extended command:

start               stop                status
services            version             reload
trace               spawn               dbsnmp_start
dbsnmp_stop         dbsnmp_status       change_password
quit                exit                set*
show*
```

SET

This command lists the operations that can be set using the SET command.

```
LSNRCTL> set
The following operations are available after set
An asterisk (*) denotes a modifier or extended command:

password            trc_file            trc_directory
trc_level           log_file            log_directory
log_status          current_listener    connect_timeout
startup_waittime    use_plugandplay
```


ORCL01808692

SHOW

This command lists the operations that can be set using the SHOW command.

```
LSNRCTL> show
The following operations are available after show
An asterisk (*) denotes a modifier or extended command:

trc_file              trc_directory         trc_level
log_file              log_directory         log_status
current_listener      connect_timeout       startup_waittime
snmp_visible          use_ckpfile           use_plugandplay
```

LSNRCTL Commands to Control the Listener

The following commands act on the listener. If there is more than one listener on the node, the command acts against the default LISTENER, unless you provide another listener name on the command line.

CHANGE_PASSWORD [*listener_name*]

You can dynamically change the password of a listener using LSNRCTL. From within LSNRCTL, enter the following:

```
change_password
```

The control utility prompts you for your old password, then for the new one. It asks you to re-enter the new one, then changes it.

> **Note:** Neither the old nor the new password displays during this procedure.

The following is an example of changing a password:

```
iris> lsnrctl

LSNRCTL for SunOS: Version 2.3.2.0.0 - on 10-Sep-95 18:59:34
Copyright (c) Oracle Corporation 1994. All rights reserved.
Welcome to LSNRCTL, type "help" for information.
LSNRCTL> change_password
Old password:
New password:
Reenter new password:
Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=iris))
Password changed for LISTENER
The command completed successfully
LSNRCTL>
```


ORCL01808693

SPAWN

Use the SPAWN command to start a program stored on the machine on which the listener runs, and which is listed with an alias in the LISTENER.ORA file.

For example, you might create a program called NSTEST, and manually add the following to the LISTENER.ORA file:

```
nstest =
(PROGRAM=(NAME=nstest)(ARGS=test1)(ENVS='ORACLE_HOME=/usr/oracle')
)
```

To run the NSTEST program, you would enter the following:

```
lsnrctl spawn nstest
```

START [*listener_name*]

Starts the named listener. If no listener name is entered, LISTENER is started by default.

STOP [*listener_name*]

Stops the named listener. If no listener name is entered, LISTENER is stopped by default.

> **Note:** You must have set a valid password if one is listed in the LISTENER.ORA parameter PASSWORDS_*listener_name* to be able to use this command.

STATUS [*listener_name*]

Displays basic information: version, start time, uptime, what LISTENER.ORA file is used, and whether tracing is turned on.

SERVICES [*listener_name*]

Provides detailed information about the services the listener listens for: how many connections have been established, how many refused. It displays three different types of services: dedicated servers from LISTENER.ORA, dispatcher information, and prespawned shadows.

> **Note:** You must have set a valid password if one is listed in the LISTENER.ORA paramete,r PASSWORDS_*listener_name*, to be able to use this command.

ORCL01808694

RELOAD [*listener_name*]

Shuts down everything except listener addresses, and re–reads the LISTENER.ORA file. This command enables you to add or change services without rebooting the system.

Note: You must have set a valid password, if one is listed in the LISTENER.ORA file parameter PASSWORDS_*listener_name*, to be able to use this command.

TRACE [*listener_name*] *level*

Turns on tracing for the listener. Choices of level are OFF, USER, or ADMIN. USER provides a limited level of tracing; ADMIN provides a more detailed trace. This command overrides the setting in the LISTENER.ORA file. (This command has the same functionality as SET TRC_LEVEL.)

Note: You must have set a valid password, if one is listed in the LISTENER.ORA file parameter PASSWORDS_*listener_name*, to be able to use this command.



Diagnostics

For detailed information on how to use tracing, see the *Oracle Network Products Troubleshooting Guide.*

VERSION [*listener_name*]

Use this command if you are asked to provide version information to Oracle Worldwide Customer Support. It provides version numbers for the following:

- listener executable version number
- TNS version number
- protocol adapter version numbers

DBSNMP_START

Use this command to start the SNMP subagent for an Oracle database running on the same node.

DBSNMP START must be run locally—you cannot run it remotely.

DBSNMP_STOP

Use this command to stop the SNMP subagent for an Oracle database running on the same node.

DBSNMP STOP must be run locally—you cannot run it remotely.

ORCL01808695

DBSNMP_STATUS

Use this command to verify whether the SNMP subagent for an Oracle database is running.

DBSNMP STATUS must be run on the same node the Oracle database is on.

SET PASSWORD *[listener_name]*

Enter this command if you want to perform administrator–only tasks on the listener. For example, you must enter the SET PASSWORD command before you can stop the listener. The password should match one listed in the LISTENER.ORA file. You may enter this command when you start up the shell or any time during your session.

The preferred, secure way to enter your password is in interactive mode. Enter the command from within LSNRCTL, for example,

```
LSNRCTL> SET PASSWORD
```

The Listener Control Utility responds:

```
enter listener password:
```

When you enter your password and press [Return], the password is not echoed on the terminal. You receive the message:

```
Command successful
```

Note: You must enter the SET PASSWORD command before you can stop the listener (with the STOP *[listener_name)*

Note: The listener supports encrypted and unencrypted passwords.

SET TRC_LEVEL [*listener_name*] *level*

Turns on tracing for the listener. Choices of level are OFF, USER, or ADMIN. Selecting USER provides a limited level of tracing; ADMIN provides a more detailed trace. This command overrides the setting in the LISTENER.ORA file. (This command has the same functionality as TRACE.)

Note: You must have set a valid password, if one is listed in the LISTENER.ORA file parameter PASSWORDS_*listener_name* to be able to use this command.



Diagnostics

For detailed information on how to use tracing, see the *Oracle Network Products Troubleshooting Guide.*

ORCL01808696

SET TRC_FILE [*listener_name*]

Use this command to set a non–default name for the trace file.

For example, the following command sets the name of the file that contains listener trace information:

```
LSNRCTL> set trc_file list.trc
```

The computer output would be something like the following:

```
Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=iris))
LISTENER parameter "trc_file" set to list.trc
The command completed successfully
```

SET TRC_DIRECTORY [*listener_name*]

Use this command to set a non–default location for the trace file or to return the location to the default.

For example, the following command sets the directory in which the trace file is placed:

```
LSNRCTL> set trc_directory /usr/oracle/admin
```

The computer output would be something like the following:

```
Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=iris))
LISTENER parameter "trc_directory" set to /usr/oracle/admin
The command completed successfully
```

SET LOG_STATUS

Logging for a listener is always ON. This command has no effect.

SET LOG_FILE [*listener_name*]

Use this command to set a non–default name for the log file.

For example, the following command sets the name of the file that contains listener log information:

```
LSNRCTL> set log_file list.trc
```

The computer output would be something like the following:

```
Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=iris))
LISTENER parameter "log_file" set to list.log
The command completed successfully
```


ORCL01808697

SET LOG_DIRECTORY [*listener_name*]

Use this command to set a non–default location for the log file or to return the location to the default.

For example, the following command sets the directory in which the log file is placed:

```
LSNRCTL> set log_directory /usr/oracle/admin
```

The computer output would be something like the following:

```
Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=iris))
LISTENER parameter "log_directory" set to /usr/oracle/admin
The command completed successfully
```

SET CURRENT_LISTENER [*listener_name*]

If there is more than one listener on a node, any LSNRCTL command acts on the default listener (LISTENER) unless another listener has been set.

For example, suppose there were two listeners on a node, LISTENER and LSNR1. If you wanted to set or show parameters for LSNR1, you would first need to send the following command from within LSNRCTL:

```
LSNRCTL> set current_listener lsnr1
```

Any subsequent LSNRCTL commands within the same LSNRCTL session would then apply to LSNR1, unless CURRENT_LISTENER were reset. For example, if the current listener had been set to LSNR1, then the STAT command would produce something like the following output:

```
LSNRCTL> stat
Connecting to (ADDRESS=(PROTOCOL=IPC)(KEY=IRIS))
STATUS of the LISTENER
------------------------
Alias                     lsnr1
Version                   TNSLSNR for SunOS: Version 2.3.1.1.0- Beta
Start Date                18-Aug-95 11:25:45
Uptime                    0 days 0 hr. 0 min. 3 sec
Trace Level               admin
Security                  OFF
SNMP                      ON
Listener Parameter File   /etc/oracle/network/admin/listener.ora
Listenr Log File          /etc/oracle/network/log/lsnr1.log
Listener Trace File       /etc/oracle/network/trace/lsnr1.trc
Services Summary...
```


ORCL01808698

```
  db1       has 1 service handler(s)
The command completed successfully
```

You can also display the current listener by using the LSNRCTL SHOW command.

> **Note:** You must enter SET CURRENT_LISTENER from within the LSNRCTL utility. When you exit the utility, the setting will be lost.

SET CONNECT_TIMEOUT TIME [*listener_name*]

This command determines the amount of time the listener will wait for a valid connection request after a connection has been started

```
LSNRCTL> set connect_timeout 20
```

The computer output would be something like the following:

```
Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=iris))
LISTENER parameter "connect_timeout" set to 20
The command completed successfully
```

SET STARTUP_WAITTIME TIME [*listener_name*]

This command sets the amount of time the listener sleeps before responding to a STATUS command:

```
lSNRCTL> set startup_waittime 10
```

The computer output would be something like the following:

```
Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=iris))
LISTENER parameter "startup_waittime" set to 10
The command completed successfully
```

SET USE_PLUGANDPLAY [*listener_name*] ON | OFF

Use this command to determine whether the Dynamic Discovery Option of Oracle Names is enabled on a listener:

```
LSNRCTL> set use_plugandplay ON|OFF
```

By default, the value of this parameter is OFF. If Oracle Names version 2 has been installed, and you want the listener to use the Dynamic Discovery Option (that is, to register itself with a well–known Names Server, set it to ON. You can also set the value in LISTENER.ORA through Oracle Network Manager, which places a slightly different parameter into LISTENER.ORA, as follows:

```
use_plugandplay_listener_name=[ON|OFF]
```


ORCL01808699

If the Dynamic Discovery Option is enabled, you can use the LSNRCTL STATUS command to see whether a service has registered itself. For example:

```
LSNRCTL> stat lsnr1
Connecting to (ADDRESS=(PROTOCOL=IPC)(KEY=IRIS))
STATUS of the LISTENER
------------------------
Alias                     lsnr1
Version                   TNSLSNR for SunOS: Version 2.3.1.1.0- Beta
Start Date                18-Aug-95 11:25:45
Uptime                    0 days 0 hr. 1 min. 49 sec
Trace Level               admin
Security                  OFF
SNMP                      ON
Listener Parameter File /etc/oracle/network/admin/listener.ora
Listener Log File         /etc/oracle/network/log/lsnr1.log
Listener Trace File       /etc/oracle/network/trace/lsnr1.trc
Services Summary...
  db1 (Registered)        has 1 service handler(s)
The command completed successfully
```

SHOW [*listener_name*] *subcommand*

All of the SET commands listed except SET PASSWORD have equivalent SHOW commands. In response to one of the SHOW comands, LSNRCTL displays the current setting of the listener for that parameter.

In addition, there are two other parameters that can be shown, but not set, through LSNRCTL.

SHOW [*listener_name*] SNMP_VISIBLE

This command displays whether the listener is accessible to SNMP clients:

```
LSNRCTL> show snmp_visible
```

The computer output would be something like the following:

```
Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=iris))
LISTENER parameter "snmp_visible" set to ON
The command completed successfully
```

Note: The SNMP_VISIBLE parameter can be displayed, but not set, through LSNRCTL:

SHOW [*listener_name*] USE_CKPFILE

Use LSNRCTL to see whether the the use of a checkpoint file has been enabled in LISTENER.ORA:

```
LSNRCTL> show use_ckpfile
```

The computer output would be something like the following:

```
Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=iris))
LISTENER parameter "use_ckpfile" set to OFF
The command completed successfully
```

If USE_CKPFILE_*listener_name* is not set in LISTENER.ORA, and you use LSNRCTL SET commands to change listener parameters, when you stop the listener, LSNRCTL sends a message reminding you that the changes you have made are not persistent.

> **Note:** You cannot set this parameter using LSNRCTL; it must be set in LISTENER.ORA. For more information about the use of checkpoint files, see "Persistent Changes" later in this chapter.

QUIT

Use this command to quit LSNRCTL and return to the operating system prompt. (Same as EXIT.)

EXIT

Use this command tp quit LSNRCTL and return to the operating system prompt. (Same as QUIT.)

Making Changes Persistent

If set to ON, the following parameter in LISTENER.ORA will make the changes you make by using the SET command in LSNRCTL persistent; that is, if you shut down the listener and then start it up again, the parameters will retain the values you set through LSNRCTL.

```
use_ckpfile_listener_name=ON
```

If this parameter is in the LISTENER.ORA file, the parameters you set using LSNRCTL are saved to a checkpoint file, *listener_name*.CKP, either in the TNS_ADMIN directory if one exists, or in the *ORACLE_HOME*/NETWORK/ADMIN directory. The values in this checkpoint file will override the values set in LISTENER.ORA.

If the USE_CKPFILE parameter is set to OFF in LISTENER.ORA, the values set by LSNRCTL are lost when the listener is restarted or the LISTENER.ORA parameters are reloaded. The default is for USE_CKPFILE to be OFF.

Note: If checkpointing is enabled and parameters in LISTENER.ORA are changed, the changes will not be visible because they are overridden by the parameters in the checkpoint file. If you want the changes in LISTENER.ORA to be visible, set the value of USE_CKPFILE to OFF and use the LSNRCTL RELOAD command to reread the LISTENER.ORA file.

Examples of the Use of LSNRCTL

This section provides some examples of the most common uses of the LSNRCTL utility.

Starting the Listener

To start a default listener, that is, a listener defined in the LISTENER.ORA file using the name LISTENER, use the command:

```
LSNRCTL START
```

Alternatively, you can enter LSNRCTL on the command line and then enter START from the program prompt.

To start a listener configured in the LISTENER.ORA file with a name other than LISTENER, include that name. For example, if the listener name is TCP_LSNR, enter:

```
LSNRCTL START TCP_LSNR
```

Or, from the LSNRCTL program prompt, just enter:

```
LSNRCTL> START TCP_LSNR
```

Stopping the Listener

To stop the default listener, use the command:

```
LSNRCTL STOP
```

To stop a running listener defined in LISTENER .ORA as TCP_LSNR, use the command:

```
LSNRCTL STOP TCP_LSNR
```

Remember, if there are any passwords in the LISTENER.ORA file, you must use the SET PASSWORD command before you can use the STOP command. You must set the password from within the LSNRCTL program; you cannot set it from the operating system command line. The method for setting the password depends on whether you are using the encrypted password feature. If you are *not* using an encrypted password, enter the password on the LSNRCTL command line. For example, the following commands stop the TCP_LSNR using an unencrypted password:

```
LSNRCTL
LSNRCTL> SET PASSWORD password
LSNRCTL> STOP TCP_LSNR
```


ORCL01808702

If you are using an encrypted password, enter the password in interactive mode. For example, the following commands stop the listener named TCP_LSNR:

```
LSNRCTL
LSNRCTL> SET PASSWORD
    Enter listener password (password is not displayed)
    Command successful
LSNRCTL> STOP TCP_LSNR
```

Stopping a listener when in batch mode is not recommended, because to do so you must include your password in a cleartext batch file, which would threaten your security. However, if you are not using an encrypted password, stopping a listener can be done by redirecting input into the command interpreter.

> **Note:** You should not stop the listener in batch mode if it requires an encrypted password.

Different operating systems use different syntax. An example for VMS and two alternative methods for UNIX follow.

To stop a listener in batch mode on VMS, create a DCL script, with a name like LSNRSTOP.COM, as follows:

```
$ lsnrctl
set password password
stop listener_name
exit
```

When you want to stop the listener, run the script as follows:

```
@LSNRSTOP
```

On a UNIX system or on OS/2, the following procedure would be effective:

Create a file with a name like LSNRSTOP that contains the following lines:

```
set password password
stop listener_name
```

You can then stop the listener from the command line by entering:

```
lsnrctl < LSNRSTOP
```

Alternatively, in UNIX you can stop the listener by creating a shell script named something like "lsnrstop". The shell script would look something like this: :

ORCL01808703

```
lsnrctl <<!
set password password
stop listener_name
exit
!
```

You can then stop the listener from the command line by entering:

```
LSNRSTOP
```

Again, Oracle Corporation recommends against stopping a listener in batch mode because of the need to expose your password.

> **Note:** Be careful when stopping a listener. On some platforms and with some protocols, when a listener is stopped any SQL*Net connections currently running are shut down. In some situations the connections continue, but it is then not possible to start the listener again until the running processes have been closed. It is good practice to send a warning message to all network users before stopping a listener.

Checking Listener Status

One of the most useful commands available with the Listener Control Utility is the STATUS command. With the status command, an administrator can get a view of the "state" of a listener by checking:



Diagnostics

- the current setting of the logging and tracing options. These are described in detail in the *Oracle Network Products Troubleshooting Guide*.
- the list of database SIDs available through this listener. These are defined in the SID mapping in LISTENER.ORA.
- whether a password is encrypted in LISTENER.ORA. (If you encrypt the listener password you can have only one password.)
- whether the network listener can respond to queries from an SNMP–based network management system
- the address the TNSLSNR is listening on

Suppose that a user wanted to know the status of a listener on a UNIX system defined in LISTENER.ORA with the default name LISTENER. The user would enter the following command at the operating system command line:

```
LSNRCTL STATUS
```

Following is sample output from a status report of a UNIX listener :

```
LSNRCTL for SunOS: Version 2.2.2.0 - Production on 15-FEB-95
07:07:10
Copyright (c) Oracle Corporation 1995. All rights reserved.
```


ORCL01808704

```
Connecting to (ADDRESS=(PROTOCOL=IPC)(HOST=orchid)(port=1334))
STATUS of the LISTENER
------------------------
Alias                     LISTENER
Version                   TNSLSNR for SunOS: Version 2.2.2.0 -
  Production
Start Date                10-FEB-94 07:06:34
Uptime                    0 days 0 hr. 0 min. 44 sec
Trace Level               ADMIN
Security                  ON
SNMP                      ON
Listener Parameter File
/private1/dvl/7012/network/admin/listener.ora
Listener Log File
/private1/dvl/7012/network/log/listener.log
Listener Trace File
/private1/dvl/7012/network/trace/listener.trc
Listening on              (ADDRESS=(PROTOCOL=tcp)(HOST=orchid)
  (port=1334)))
                          (ADDRESS=(PROTOCOL=decnet)(node=23.106)
  (object=orchid)))
Services Summary...
  orchid           has 1 service handlers
The command completed successfully
```

Retrieving Listener Services

A database administrator who wanted to get information about the services of the listener would enter the following commands from within LSNRCTL:

```
LSNRCTL>[SET PASSWORD password]
LSNRCTL> SERVICES
```

The output of a LSNRCTL SERVICES command follows:

```
LSNRCTL for SunOS: Version 2.1.3.0.0 - Production on 10-FEB-94
07:14:55

Copyright (c) Oracle Corporation 1993.  All rights reserved.

Connecting to (ADDRESS=(PROTOCOL=IPC)(KEY=ruth))
Services Summary...
  ruth             has 1 service handlers
    DEDICATED SERVER established:99 refused:0
The LSNRCTL command completed successfully
```

In this example, the LSNRCTL SERVICES command returned the information that the listener had established 99 connections using a dedicated server process and refused none.

ORCL01808705

In the following example the LSNRCTL SERVICES command returns information about four types of service handlers.

```
LSNRCTL for SunOS: Version 2.2.2.0.0 - Production on 15-Mar-1995 17:41:12

Copyright (c) Oracle Corporation 1994.  All rights reserved.

Connecting to (ADDRESS=(PROTOCOL=ipc)(KEY=orchid))
Services Summary...
   listener          has 4 service handlers
    DEDICATED SERVER established:0 refused:0
    PRESPAWNED SERVER established:0 refused:0 current:0 max:1 state:ready
      PID:15439
      (ADDRESS=(PROTOCOL=ipc)(DEV=4)(KEY=#15439.1))
    PRESPAWNED SERVER established:5 refused:0 current:0 max:1 state:ready
      PID:15441
      (ADDRESS=(PROTOCOL=tcp)(DEV=4)(HOST=139.185.22.25)(PORT=3334))
    DISPATCHER established:30 refused:0 current:7 max:21 state:ready
      D000 (machine: orchid, pid: 15406)
      (ADDRESS=(PROTOCOL=tcp)(DEV=7)(HOST=139.185.22.25)(PORT=3330))
The command completed successfully
```

This message shows that since this listener has been running, it has made a total of five connections using a prespawned server for TCP/IP, and that it has one running at this time. It also shows that there are seven connections currently established using a multi-threaded server dispatcher, with a total of 30 established since the process started.

Modifying Client Parameters

Client configuration parameters are contained in the SQLNET.ORA file. Generally, you generate SQLNET.ORA files for clients with similar navigation needs (called *client profiles* or *client types*) using Oracle Network Manager. The SQLNET.ORA file determines the following client characteristics:

- preferred choices of name respolution services
- preferred Names Servers (if Oracle Names is used on the network)
- preferred Oracle MultiProtocol Interchanges (if used on the network)
- non–default tracing and logging characteristics
- preferred security services (if Secure Network Services is used on the network)

ORCL01808706

- characteristics of connection requests (whether to use a dedicated server, out-of-band breaks, IPC addresses)
- client registration information for statistical reporting

The SQLNET.ORA file created by Network Manager generally is used by a number of clients. If you want to change any of these parameters for a whole group of clients, you may use Network Manager to change the network definition and create new SQLNET.ORA files to be distributed as needed.

Alternatively, if you want to change parameters for individual clients, you may use the SQLNET.ORA Editor. This utility is part of the Client Status Monitor. It enables you to open the SQLNET.ORA file on an individual client and change its parameters either through a graphical user interface or through the command line.

If you are using Oracle Names version 2.0 and the Dynamic Discovery Option, you do not necessarily need to use Network Manager. If you do not use it, a SQLNET.ORA file is not created. If you find you want to provide some non–default parameters for the client, you can create an empty SQLNET.ORA using any text editor, such as vi, and then add the parameters you need using the SQLNET.ORA Editor. Using this method to create the file makes it less likely that typographical or formatting mistakes will cause errors.

Controlling Network Services from Network Manager

Oracle Network Manager version 3.0 lets you run some Listener Control Utility, Names Control Utility, and Interchange Control Utility commands without having to go to a command line prompt to execute them.

After your network has been defined, you can obtain information on any network service by selecting it on the Map View or Tree View Object List and then selecting Control Network Service from the Special menu.

> **Note:** You must define the network configuration and distribute the configuration files throughout the network before you can use this option.

A control window appears. Select a command from the drop-down list. (The commands available depend on the type of network object you have selected.) For example, if you select Trace, select the level of tracing you want from the Level list. The result of your command will be displayed in the large read–only box on this window.

ORCL01808707

Following are the commands you can use from Network Manager to control network services remotely.

LSNRCTL	VERSION STATUS RELOAD SET_TRACE_LEVEL SERVICES
NAMESCTL	VERSION STATUS PING SET TRACE_LEVEL SHOW TRACE_LEVEL
INTCTL	VERSION STATUS



ONames



Intchg

The Listener Control Utility (LSNRCTL) commands are described in "Using the Listener Control Utility" in this chapter. See the *Oracle Names Administrator's Guide* for information on NAMESCTL commands. See the *Oracle MultiProtocol Interchange Administrator's Guide* for information on INTCTL commands.

Initiating a SQL*Net Connection

There are a number of ways to initiate a connection with an Oracle server. Commonly used methods are:

- the operating system command line
- a tool logon screen
- a 3GL program
- special commands within certain tools

The specifics of use are slightly different in each case. Each of the general methods listed is briefly covered here. To identify the method used in a specific tool, refer to the tool's user's guide.

Connecting from the Operating System Command Line

The general form of connecting an application to a database server from the command line is:

```
tool username/password@service_name
```

In this syntax:

ORCL01808708

tool Specifies the command used to invoke a tool such as SQL*Plus, SQL*Forms, etc.

username Specifies an Oracle username on the server.

password Specifies the corresponding password on the server.

service_name Specifies a service name entered in the TNSNAMES.ORA file that identifies the connect descriptor for the desired server. If the server is in the client's default domain, the service name does not need to include the domain name. However, if the server is in another domain, the service name must include the domain. (The default domain is determined by a parameter in the client's SQLNET.ORA file. See the section on the SQLNET.ORA file in Appendix A of this manual.)

For example, in a network with only one domain, the default .WORLD domain, it is not necessary to include .WORLD in the service name. For example:

```
% sqlplus scott/tiger@SERVERX
```

However, if the client's default domain were .EAST and the server's domain were .WEST, then the service name would have to include the domain. For example,

```
% sqlplus scott/tiger@SERVERX.WEST
```

Note: To prevent the password from displaying during a logon, you can leave out the password parameter on the command line; you will then be prompted to enter your password without it showing on screen.

Most Oracle tools can use the operating system command line to connect; some provide alternatives.

Connecting from the Tool Logon Screen

Some tools provide a logon screen as an alternative form of logon. A user can log on to a database server just as easily by identifying both the username and service name in the username field of the tool logon screen, and typing the password as usual in the password field. Figure 5 – 1 shows a SQL*Forms logon screen where the user SCOTT is connecting to the server SERVERX with a password of TIGER. Notice the password cannot be seen, a standard feature of Oracle tool logon screens.

ORCL01808709

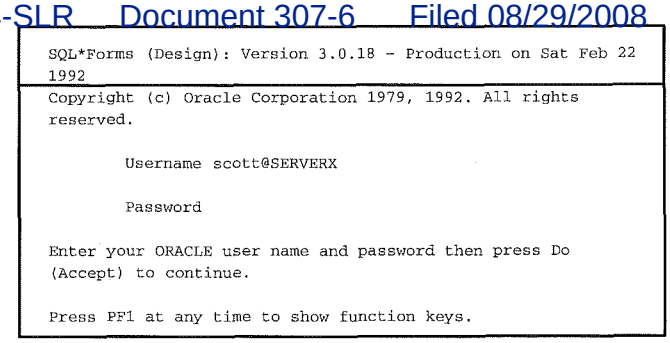


Figure 5 – 1 Connection from Logon Screen

Connecting from a 3GL Application

In applications written using a 3GL, the program must establish a connection to a server using the following syntax:

```
EXEC SQL CONNECT :username IDENTIFIED BY :password
```

In this connection request, the *:username* and *:password* are 3GL variables that can be set within the program either statically or by prompting the user. When connecting to a database server, the value of the *:username* variable is in the form:

```
username@service_name
```

which is the same as in the tool logon screen above. The *:password* variable contains the password for the database account being connected to.

Connecting Using Special Commands within Tools

Some Oracle tools have commands for database connection, once the tool has been started, to allow an alternative username to be specified without leaving the tool. Both SQL*Plus and SQL*DBA allow the CONNECT command using the following syntax:

```
SQL> CONNECT username/password@service_name
```

For example:

```
SQL> CONNECT SCOTT/TIGER@SERVERX
```

This is very similar to the operating system command line method, except that it is entered in response to the tool prompt instead of the operating system prompt.

Other Oracle tools use slightly different methods specific to their function or interface. For example, Oracle CDE tools use logon buttons and a pop–up window with the username, password, and remote database ID field. For more information on connecting to Oracle with a specific tool, refer to the tool's user guide.

Distributed Database Management

SQL*Net plays a major role in distributed database management. In a distributed transaction, SQL*Net provides the means for clients and servers to communicate by way of their SQL–based dialog language. SQL*Net performs transparently to enable distributed database functions. This section highlights the database functionality used in conjunction with SQL*Net.

Database Links

A *database link* is a database object that links an Oracle account in one database to an Oracle account in another database. The data in the remote schema is then accessible to the database in which the database link is defined.

> **Note:** If your network uses Oracle Names, a global database link is created on every database on the network to every other database. Therefore, you do not need to create additional database links, as described in this section. See "Database Links wth Oracle Names" later in this chapter.

The generic syntax for creating a database link in SQL is:

```
CREATE [PUBLIC] DATABASE LINK linkname
[CONNECT TO username IDENTIFIED BY password]
USING 'service_name'
```

In this syntax:

[PUBLIC]	Specifies a database link available to all users with the CREATE SESSION privilege. If the PUBLIC option is omitted, a private link available only to the creator is created. Note that creating a public database link requires CREATE PUBLIC DATABASE LINK privilege.
linkname	Specifies the name of the database link. If the remote server is in the local server's domain, the link name does not need to include the domain name. However, if the server is in another domain, the link name must include the domain. (The domain is determined by DB_DOMAIN in the initialization parameter file).
CONNECT TO	Optionally specifies a single username and password for all users of the database link to share. If the clause is omitted, the Oracle username and password of the user account using the database link will be used to connect to the remote database server.

ORCL01808711

username	Specifies a valid Oracle username on the remote database server.
password	Specifies the corresponding password of the username on the remote database server.
service_name	Specifies the service name defined in the TNSNAMES.ORA file or stored in Oracle Names associated with the connect descriptor for the desired database. If the remote server is in the local server's default domain, the service name does not need to include the domain name. However, if the server is in another domain, the service name must include the domain. (The default domain is determined by a parameter in the server's SQLNET.ORA file. See the section on the SQLNET.ORA file in Appendix A of this manual.)

Prior to Oracle7, a database administrator could specify any linkname for a database link. However, with Oracle7 and later, a database link must have the same name as the global database name of the database. Remember that the service name is also the same as the global database name; therefore, the linkname and service name are now the same.

For example, the command for creating a public database link to a database which has the global database name ORCHID.HQ.ACME is as follows:

```
CREATE PUBLIC DATABASE LINK ORCHID.HQ.ACME
CONNECT TO scott IDENTIFIED BY tiger
USING 'ORCHID.HQ.ACME'
```

> **Note:** The CONNECT TO *username* IDENTIFIED BY *password* clause and the USING '*global_database_name*' clause are both optional.

Public Database Links with a Default Connection

Figure 5 – 2 shows a public database link created by the DBA user SYSTEM using the service name NY_FIN.HQ.ACME. The link is created by entering:

```
CREATE PUBLIC DATABASE LINK NY_FIN.HQ.ACME
USING 'NY_FIN.HQ.ACME'
```


ORCL01808712

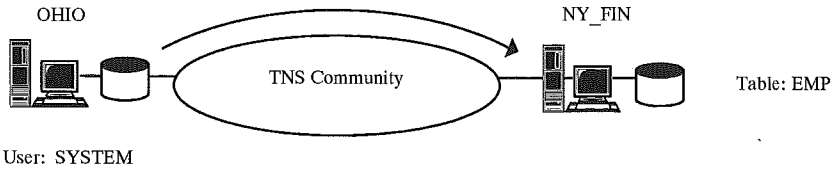


Figure 5 – 2 Public Database Link with Default Connection

Users connected to OHIO.SALES.ACME can use the NY_FIN.HQ.ACME database link to connect to NY_FIN.HQ.ACME with the same username and password they have on OHIO.SALES.ACME. To access the table on NY_FIN.HQ.ACME called EMP, any user could issue the SQL query:

```
SQL> SELECT * FROM EMP@NY_FIN.HQ.ACME;
```

> **Note:** If the target database were in the default domain of the source database, the user would not need to include the domain in the link name or service name, or in the SELECT command.

This query would initiate a connection from OHIO to NY_FIN using the current username and password to log onto NY_FIN. The query would then be processed on NY_FIN, and the data available to the current user from the table EMP would be returned to OHIO. Each user creates a separate connection to the server. Subsequent queries to that database link by that user would not require an additional logon.

Public Database Links with a Specific Connection

Figure 5 – 3 shows the database link created by the user SYSTEM with the service name NY_FIN:

```
CREATE PUBLIC DATABASE LINK NY_FIN
CONNECT TO FINPUBLIC IDENTIFIED BY NOPASS
USING 'NY_FIN'
```

> **Note:** The CONNECT TO *username* IDENTIFIED BY *password* clause and the USING *'global_database_name'* clause are both optional.

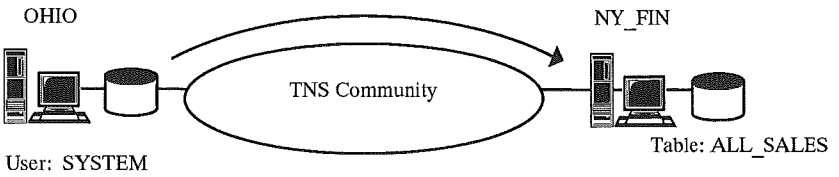


Figure 5 – 3 Public Database Link with Specific Connection

ORCL01808713

Any user connected to OHIO can use the NY_FIN database link to connect to NY_FIN with the common username/password of FINPUBLIC/NOPASS. To access the table in the FINPUBLIC account of NY_FIN called ALL_SALES, any user could issue the SQL query:

```
SQL> SELECT * FROM ALL_SALES@NY_FIN;
```

This query would initiate a connection from OHIO to NY_FIN to the common account FINPUBLIC. The query would be processed on NY_FIN and data from the table ALL_SALES would be returned to OHIO.

Each user creates a separate connection to the common account on the server. Subsequent queries to that database link by that user would not require an additional logon.

Connection Qualifiers

You can also define *connection qualifiers* to database links. Connection qualifiers provide a way to create more than one link to a given database. Alternate links are a useful way to access different accounts on the same database with different sets of access privileges. The alternate link created by a connection qualifier must include a reference to a database by its global database name (or service name).

A connection qualifier contains a qualifier name and, optionally, a username and password. To create a connection qualifier, use a statement similar to the following:

```
CREATE PUBLIC DATABASE LINK NY_FIN@PROFITS
CONNECT TO ACCOUNTS IDENTIFIED BY TAXES
USING 'NY_FIN'
```

To use the connection qualifier, you append the qualifier name to the service name of the database you want to access.

For example, the following SQL queries use three separate database links to the same database, using different connection qualifiers:

```
SELECT * FROM EMP@NY_FIN;
SELECT * FROM SCHEDULE@NY_FIN@PROFITS;
SELECT * FROM EMPSALARIES@NY_FIN@FIN;
```

In this example @PROFITS and @FIN are connection qualifiers.

ORCL01808714

Dropping a Database Link

You can drop a database link just as you can drop a table or view. The command syntax is:

```
DROP DATABASE LINK linkname;
```

For example, to drop the database link NY_FIN, the command would be:

```
DROP DATABASE LINK NY_FIN;
```

Finding Available Database Links



Server

Any user can query the data dictionary to determine what database links are available to that user. For information on viewing the data dictionary, refer to the *Oracle7 Server Concepts* or the *Oracle7 Administrator's Guide.*

Database Links with Oracle Names

When you define a network that includes Oracle Names, Network Manager automatically creates a global database link to every database server you define from every other database server in the network. These database links do not reside in the data dictionary, but in the network definition to which the Names Servers refer. The database links thus created do not initially include a CONNECT TO clause, so that users reach the linked database using the same usernames and passwords as they use to reach the first database.

```
SQL> SELECT * FROM EMP@OHIO, DEPT@NY_FIN;
```

Explicitly Defined Database Links



NetMan

You can edit global database links to include CONNECT TO data using Network Manager. When you edit a database, you can specify a single default username and password for the database link. See Chapter 5 in the *Oracle Network Manager Administrator's Guide* for details on how to edit database links.

Connection Qualifiers



NetMan

You can also define *connection qualifiers* to global database links through Network Manager. Connection qualifiers provide a way to create multiple links to the same database. Multiple database links to the same database provide different access routes with different accounts and privileges. See Chapter 5 in the *Oracle Network Manager Administrator's Guide* for details on how to create connection qualifiers using the Network Manager.



ONames

For a more detailed discussion of database links, see *Oracle7 Server Distributed Systems, Volume I,* and the *Oracle Names Administrator's Guide.*

ORCL01808715

Synonyms

Database *synonyms* are a standard SQL feature used to provide alternate names for database objects and, optionally, their locations. A synonym can be created for any table, view, snapshot, sequence, procedure, function, or package. All synonyms are stored in the data dictionary of the database in which they are created. To simplify remote table access through database links, a synonym can be defined to allow single–word access to remote data, isolating the specific object name and the location from users of the synonym. The syntax to create a synonym is:

```
CREATE [PUBLIC] SYNONYM_name
FOR [schema.]object_name[@database_link_name]
```

In this syntax:

[PUBLIC]	Specifies that this synonym is available to all users. Omitting this parameter makes a synonym private, and usable only by the creator. Public synonyms can be created only by a user with CREATE PUBLIC SYNONYM system privilege.
synonym_name	Specifies the alternate object name to be referenced by users and applications.
schema	Specifies the schema of the object specified in object_name. Omitting this parameter uses the creator's schema as the schema of the object.
object_name	Specifies either a table, view, sequence, or other name as appropriate.
database_link_name	Specifies the database link which identifies the remote username in which the object specified in object_name is located.

A synonym must be a uniquely named object for its schema. If a schema contains a database object and a public synonym exists with the same name, Oracle always finds the database object when the user that owns the schema references that name.

Because a synonym is merely a reference to the actual object, the security domain of the object is used when the synonym is accessed. For example, a user that has access to a synonym for a specific table must also have privileges on that table to access the data in it. If the user attempts to access the synonym, but does not have privileges on the table it identifies, an error occurs indicating that the table or view does not exist.

Figure 5 – 4 shows two servers, OHIO and NY_FIN, in which a database link from OHIO to NY_FIN and the synonym FOR_SALE provide an

ORCL01808716

alternate object name for use in OHIO to reference the OPEN table in NY_FIN. The database link and the synonym are created as follows:

```
CREATE PUBLIC DATABASE LINK NY_FIN
CONNECT TO REAL_ESTATE IDENTIFIED BY NOPASS;
USING 'NY_FIN'
CREATE PUBLIC SYNONYM FOR_SALE
FOR OPEN@NY_FIN;
```

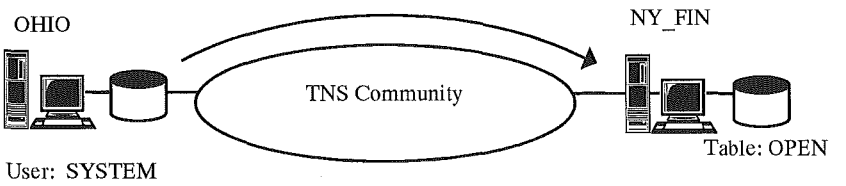


Figure 5 – 4 Using Synonyms for Alternate Object Names

The table OPEN on NY_FIN could be accessed from OHIO using the SQL statement:

```
SELECT * FROM FOR_SALE;
```

Using this database link, the user is logging on to NY_FIN as user REAL_ESTATE. Notice that this public synonym was created by the DBA on behalf of the REAL_ESTATE username. If the table OPEN were owned by another user, such as SALES_MGR, the CREATE SYNONYM statement would have referred to the object as SALES_MGR.OPEN@NY_FIN. Without such a prefix, a table that does not exist in the database link user's schema would return an error, since it would be looking for the OPEN table owned by the REAL_ESTATE user.

Maintaining Location Transparency

When using a synonym to access a database object over a database link, the user of the synonym is said to have *location transparency*. For example, an application developer using the FOR_SALE synonym from the previous example has the illusion that FOR_SALE is a database object available for use as any other object would be. The reference to the database link is invisible to the developer; therefore, any application built using the synonym would have no reference to the location of specific data.

This ability to isolate applications from the location of data in a distributed transaction ensures maximum flexibility for future enhancements or changes to the application. For example, if the OPEN table were to move from one database server to another, only the synonym or the database link would need to be changed to identify the

new location. The applications would continue to reference the same object name, although they would be connecting to a new location to access the data in that table. Figure 5 – 5 shows the most common method of redefining the location of a table to retain location transparency.

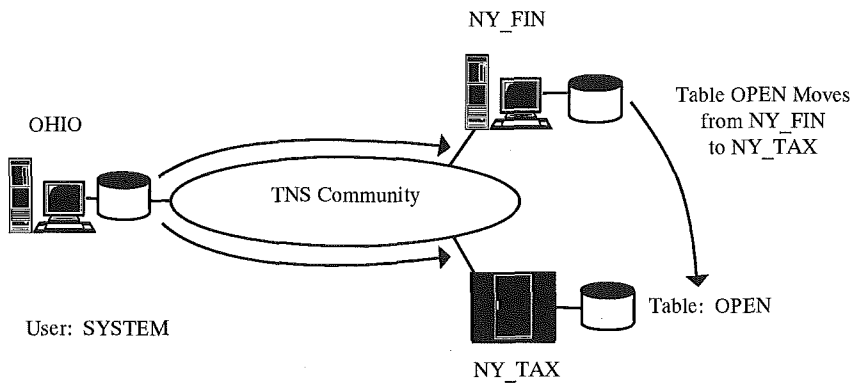


Figure 5 – 5 Redefining Table Location to Retain Location Transparency

```
CREATE PUBLIC DATABASE LINK NY_TAX
CONNECT TO REALTOR IDENTIFIED BY NOPASS
USING 'NY_TAX'
DROP SYNONYM FOR_SALE;
CREATE PUBLIC SYNONYM FOR_SALE
FOR OPEN@NY_TAX;
```

To relocate the table, a second database link was created called NY_TAX that connected to a new database with the service name NY_TAX, and the synonym was recreated to reference the NY_TAX database link instead of the NY_FIN database link. Any other tables that were accessed through the NY_FIN database link to NY_FIN would continue to function properly.

Alternatively, if the only table being accessed on NY_FIN were the OPEN table, the synonym could have remained unchanged and the database link redefined to use the service name NY_TAX instead of NY_FIN. Either option is equally effective.

Snapshots

An Oracle system with both the distributed option and the procedural option can replicate tables that are frequently queried by users on many nodes of a distributed database. By having read–only copies of heavily accessed data on several nodes, the distributed database does not need to send information across a network repeatedly, thus helping to improve the performance of the database application. Oracle provides an automatic method for table replication called *snapshots*. Snapshots are

ORCL01808718

read–only copies of a master table located on a remote node. A snapshot can be queried, but not updated; only the master table can be updated. A snapshot is periodically refreshed to reflect changes made to the master table.

Maintaining snapshots of a master table among the nodes of a distributed database is often a useful feature for the following reasons:

- Queries can be issued against a local snapshot. Therefore, associated query performance is fast because the requested data does not have to be shipped over a network.
- If the database that contains the master table or a database that contains a snapshot experiences a failure or is cut off from the distributed database by a network failure, the other (read–only) copies remain accessible to other users of the system.

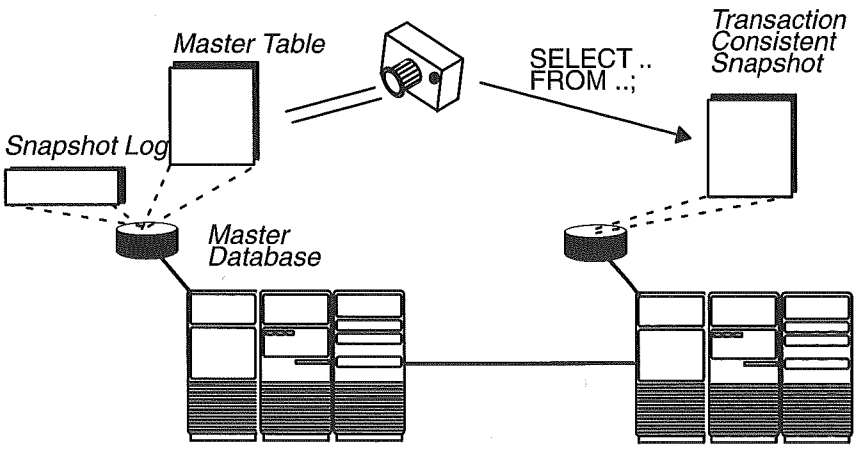


Figure 5 – 6 Table Replication Using Snapshots

A snapshot is a full copy or a subset of a table that reflects a recent state of the master table. A snapshot is defined by a distributed query that references one or more master tables, view, or other snapshots. A database that contains a master table is referred to as the master database.

Simple vs. Complex Snapshots Each row in a simple snapshot is based on a single row in a single remote table. Therefore, a simple snapshot's defining query has no GROUP BY or CONNECT BY clauses, or subqueries, joins, or set operations. If a snapshot's defining query contains any of these clauses or operations, it is referred to as a *complex snapshot*.

ORCL01808719

Creating Snapshots

Create a local snapshot using the SQL command CREATE SNAPSHOT. As when creating tables, you can specify storage characteristics for the snapshot's data blocks, extent sizes and allocation, and the tablespace to hold the snapshot; or you can specify a cluster to hold the snapshot. Unique to snapshots, you can specify how the snapshot is to be refreshed and the distributed query that defines the snapshot. You must fully qualify any remote table names used in the defining query. For example, the following CREATE SNAPSHOT statement defines a local snapshot to replicate the remote EMP table located in the SCOTT schema in NY:

```
CREATE SNAPSHOT emp_sf
   PCTFREE 5 PCTUSED 60
   TABLESPACE users
   STORAGE (INITIAL 50K NEXT 50K PCTINCREASE 50)
   REFRESH FAST
           START WITH sysdate
           NEXT sysdate + 7
   AS SELECT * FROM scott.emp@ny;
```

Whenever a snapshot is created, it is immediately populated with the rows returned by the query that defines the snapshot. Thereafter, the snapshot is refreshed as specified by the REFRESH clause; see *Oracle7 Server Distributed Systems, Volume 2: Replicated Data* for more information about refreshing snapshots.

When a snapshot is created, Oracle creates several internal objects in the schema of the snapshot. These objects should not be altered. At the snapshot node, Oracle creates a *base table* to store the rows retrieved by the snapshot's defining query. Oracle then creates a read–only view of this table that is used whenever queries are issued against the snapshot.

Specifying the Defining Query of a Snapshot

The defining query of a snapshot can be any valid query of tables, views, or other snapshots that are not owned by user SYS. The query cannot contain either an ORDER BY or FOR UPDATE clause. Furthermore, simple snapshots are defined using a query that does not contain GROUP BY or CONNECT BY clauses, or join, subquery, or set operations.

ORCL01808720

The query that defines a snapshot can define a snapshot with a different structure from that of the master table. For example, the following CREATE SNAPSHOT statement creates a local snapshot named EMP_DALLAS, with only the EMPNO, ENAME, and MGR columns of the master table (in New York), and only the rows of the employees in department 10:

```
CREATE SNAPSHOT emp_dallas
   .
   .
   .
   AS SELECT empno, ename, mgr
      FROM scott.emp@ny
      WHERE deptno = 10;
```

Privileges Required to Create Snapshots

To create a snapshot, the following sets of privileges must be granted as follows:

- To create a snapshot in your own schema, you must have the CREATE SNAPSHOT, CREATE TABLE, and CREATE VIEW (only for simple snapshots) system privileges as well as the appropriate SELECT privilege on the master tables.
- To create a snapshot in another user's schema, you must have the CREATE ANY SNAPSHOT, CREATE ANY TABLE, CREATE ANY VIEW, and CREATE ANY INDEX (only for simple snapshots) system privileges as well as the appropriate SELECT privilege on the master table. Additionally, for the snapshot to be of any use, the owner of the snapshot must have the appropriate SELECT privilege on the master tables

In both of the above cases, the owner of the snapshot must also have sufficient quota on the tablespace intended to hold the snapshot.

The large set of privileges required to create a snapshot is due to the underlying objects that must also be created on behalf of the snapshot.

Refreshes

Periodically, a snapshot is refreshed to reflect the current state of its master table. To refresh a snapshot, the snapshot's defining query is issued and its results are stored in the snapshot, replacing the previous snapshot data. Each snapshot is refreshed in a separate transaction, either automatically by Oracle (according to the interval set when the snapshot was defined or altered) or manually. Instructions for refreshing table snapshots are included in the *Oracle7 Server Administrator's Guide*.

Snapshot Logs

A simple snapshot can be refreshed from a snapshot log to expedite the refresh process. A *snapshot log* is a table in the master database that is

ORCL01808721

associated with the master table. Oracle uses a snapshot log to track the rows that have been updated in the master table when a certain simple snapshot based on the master table is refreshed. One snapshot log can be used by multiple simple snapshots). Only the appropriate rows in the snapshot log need to be applied to the snapshot to refresh it (called a *fast refresh*). If no other simple snapshot requires an applied row in the log, it is purged from the log to keep the log size small; however, if another simple snapshot requires the row for its next refresh, the row remains in the log.

A complex snapshot or simple snapshot without a snapshot log must be completely regenerated using the master tables every time the snapshot is refreshed (called a *complete refresh*).

Creating Snapshot Logs

Snapshot logs are created in the master database using the SQL command CREATE SNAPSHOT LOG. You can set storage options for the snapshot log data blocks, extent sizes and allocation, and tablespace to hold the snapshot log. The following statement creates a snapshot log associated with the EMP table:

```
CREATE SNAPSHOT LOG ON emp
   TABLESPACE users
   STORAGE (INITIAL 10K NEXT 10K PCTINCREASE 50);
```

The snapshot log is always created in the same schema that contains the master table. Since you cannot specify a name for the snapshot log (one is implicitly given by Oracle), uniqueness is not a concern.

If you own the master table, you can create an associated snapshot log if you have the CREATE TABLE and CREATE TRIGGER system privileges. If you are creating a snapshot log for a table in another user's schema, you must have the CREATE ANY TABLE and CREATE ANY TRIGGER system privileges. In either case, the owner of the snapshot log must have sufficient quota in the tablespace intended to hold the snapshot log.

The privileges required to create a snapshot log directly relate to the privileges necessary to create the underlying objects associated with a snapshot log.

Using Snapshots

Snapshots are queried just like a table or view. For example, the following statement queries a snapshot named EMP:

```
SELECT * FROM emp;
```

To query a snapshot, you must have the SELECT object privilege for the snapshot, either explicitly or via a role.

ORCL01808722

In release 7.0 of the Oracle Server, snapshots are read–only. You cannot issue any INSERT, UPDATE, or DELETE statements when using a snapshot; if you do, an error is returned. Although INSERT, UPDATE, and DELETE statements can be issued against the base table for a snapshot, they can corrupt the snapshot. Never manipulate data in the base table of a snapshot. Updates are allowed on the master table only, which must then be refreshed to update the snapshot.

Creating Views and Synonyms Based on Snapshots

Views or synonyms can be defined based on snapshots. The following statement creates a view based on the EMP snapshot:

```
CREATE VIEW sales_dept AS
   SELECT ename, empno
   FROM emp
   WHERE deptno = 10;
```

Dropping Snapshots

You can drop a snapshot independently of the master tables or a snapshot log if you decide that you no longer want to replicate a table in a database. To drop a local snapshot, use the SQL command DROP SNAPSHOT. For example:

```
DROP SNAPSHOT emp;
```

If you drop the only snapshot of a master table, you should also drop the snapshot log of the master table, if appropriate.

Only the owner of a snapshot, or a user with the DROP ANY SNAPSHOT, DROP ANY TABLE, and DROP ANY VIEW system privileges can drop a snapshot.

For more information about table replication, snapshots, and distributed database design, refer to the following manuals: *Oracle7 Server Distributed Systems, Volumes I and II*, and the *Oracle7 Server Administrator's Guide.*

Copying Data between Databases

When working with a distributed database, you may want to move data from one database to another. For example, you may want to download sales data from a particular region to your workstation for further analysis. The SQL*Plus COPY command lets you copy data from:

- a remote database to your local database
- your local database to a remote database
- one table on a remote database to another table on that database

ORCL01808723

- a table on one remote database to a table on another database

You do not need to be directly connected to any of the databases in order to copy the information from one to another. These operations, except the first, are impossible to perform with database links and queries. Using the COPY command is similar to using the CREATE DATABASE LINK command, in that you must know:

- a valid service name for the remote database
- a valid username and password for the remote database

The COPY command syntax is:

```
COPY FROM username/password@service_name
[TO username/password@service_name]
(APPEND | CREATE | INSERT | REPLACE)
tablename [columnlist]
USING subquery
```

In this syntax:

service_name	Specifies a valid service name in the TNSNAMES.ORA file. Each service name maps to a connect descriptor that connects to a database using the keyword SID.
APPEND	Specifies that if the destination table already exists, COPY will insert the copied data into it. If the table does not already exist, COPY will create it and then insert the copied data into it.
CREATE	Specifies that COPY will create a new table and insert the copied data into it. If the table already exists, COPY will report an error and stop processing.
INSERT	Specifies that if the destination table already exists, COPY will insert the data into it. If the table does not already exist, COPY will report an error and stop processing.
REPLACE	Specifies that if the table already exists, it will be dropped and replaced with the new data. If the table does not already exist, it will be created and the data inserted into it.
tablename	Specifies the destination table to which you are copying data.

ORCL01808724

columnlist	Specifies that only data from the specified columns should be copied.
subquery	Specifies the query to be used to select the data to be copied.

Note that you can use any kind of query to select the data to be copied. For example, you can COPY select rows from the source table by using a WHERE clause. You can also COPY specific columns from the source table by listing those columns after the SELECT statement, instead of using the asterisk that means all columns. The copied columns will have the same names in the destination table that they had in the source table.

Following is an example of a COPY command that copies only two columns from the source table specified by the alias BOSTON, and copies only those rows in which the value of DEPTNO is 30. Since the TO clause is omitted, the table is copied to the database to which SQL*Plus is currently connected:

```
COPY FROM SCOTT/TIGER@BOSTON
     REPLACE EMPCOPY2
     USING SELECT ENAME, SAL
          FROM EMP
          WHERE DEPTNO = 30;
```

In the following example, the command copies the entire table EMP from the remote database specified by the SALES service name in the TNSNAMES.ORA file. Because the TO clause is omitted, the table is copied to the database to which SQL*Plus is currently connected. The new table is named EMP2:

```
COPY FROM SCOTT/TIGER@SALES
     CREATE EMP2
     USING SELECT * FROM EMP;
```

You can also copy data to a local table from a remote table using the syntax:

```
COPY FROM SCOTT/TIGER@servicename
     INSERT INTO local_table
     SELECT * FROM remote_table@dblink
```

For further information about the COPY command, see the *SQL*Plus User's Guide and Reference Manual.*

ORCL01808725

Support for Operating System Authorized Logins

In SQL*Net version 1, Oracle provided support for operating system authorized logins across a network using the special database login accounts. This support continues in SQL*Net version 2 on the same platforms and protocols for compatibility of function, but is not expanded to additional platforms and protocols.

Operating System Authorized Logins

Operating system authorized logins to a database allow a user who has operating system privileges on a machine to get automatic access to an Oracle database on the machine without supplying a database password. The assumption is that having supplied a password to the operating system, that user can be trusted to access his data in his own operating–system–authorized user account on that machine. The default prefix for this account is OPS$; however, the database administrator may assign some other prefix.

Extending this notion over a network, operating system authorized logins allow users to log into a database on a remote machine using their local userid, without supplying a database password. The operating system user account must match the username of the operating system authorized account.

For further information about user authentication, see the *Oracle7 Server Distributed Systems, Volume I.*

Warning about Operating System Authorized Logins

It is not uncommon for many employees to have DOS or other types of workstations on their desks over which they have complete control. It is a trivial task for such a person to create a username of any valid string of characters. If John Doe is the president of a company and uses operating system-authorized logon accounts across the network, any user can create a JDOE account on a personal workstation and access John's data without a password. On some server platforms this is more difficult than described, but never impossible. This is a particularly serious problem where an OPS$ account has DBA privileges. In such cases, it is very easy to masquerade as an Oracle DBA.

Oracle's Recommendation Regarding Operating System Authorized Logins

Because of these security implications, Oracle Corporation recommends against using operating system authorized logins to remote databases. This does not mean you cannot use operating system authorized login accounts in your database; it simply prevents them from being used across the network. The same OPS$ account can still be accessed if a password is supplied in connections over the network.

If you have applications that rely on operating system authorized logins across a network, you should consider altering them to require the user to supply a password when connecting over the network.

In SQL*Net version 1 the ability to perform operating system authorized logins exists in some SQL*Net TCP/IP and SQL*Net DECnet drivers on UNIX and VMS systems. The default setting is for operating system authorized logins to be enabled unless explicitly disabled. We recommend that you explicitly disable operating system authorized logons where they exist in any version 1 drivers you have installed. For information regarding your platform and SQL*Net Version 1 drivers, and how operating system authorized logins can be disabled, see your Oracle operating system specific manuals regarding SQL*Net version 1.



OS Doc

In SQL*Net version 2, operating system authorized login support is enhanced to make it more difficult to act as an imposter, but it is still not impossible. By default, operating system authorized logins are disabled in SQL*Net version 2.

> **Note:** If externally authenticated logins are enabled (by setting the REMOTE_OS_AUTHENT parameter to TRUE), it will be possible to connect without a password even though the underlying SQL*Net protocol cannot verify the identity of the user. For this reason, REMOTE_OS_AUTHENT should always be set to FALSE (the default value).
>
> If it is necessary to make externally authenticated logins, then a secure protocol such as the Bequeath (Pipe) driver should be used.

In all cases with SQL*Net version 1 or version 2, Oracle recommends leaving remote DBA privileges disabled.

Support for Network Authentication Adapters in Release 2.2

With SQL*Net release 2.2 and Oracle7 Server release 7.2, the ability to use network authentication services for single sign–on and smart card authentication is provided. Use Oracle Network Manager to configure these authentication adapters. A site will have a choice of which adapters to link into their SQL*Net configuration. The Secure Network Services version 2 product is required to enable these authentication adapters. For information on configuring network authentication adapters, see the *Secure Network Services Administrator's Guide.*



SNS

A discussion of security issues in networked systems is in Chapter 6 of *Oracle7 Server Distributed Systems, Volume I.*

Authentication Services

Authentication services for single sign–on servers or smart cards are usually part of a Network Operating System (NOS) which overlays

ORCL01808727

several machines. Administration of a network with many machines can be centralized by creating "network users" who have the same identity and privileges no matter which machine they are actually using.

Network Identity

All authentication services provide the concept of a network identity so that no matter which machine a user is logged into, he or she can log onto the network and be identified as the same person. This method of authentication ensures that a user is who he or she claims to be.

Network Roles

Because most operating systems support some concept of roles, groups, or identifiers, many authentication services such as single sign–on servers based on Kerberos can support the concept of network roles. For example, an administrator creates the role and then assigns it to certain network users. Role definitions are consistent across all machines linked by the authentication service.

Secure External Authentication Logins

Using the network identity provided by the authentication service, the database can provide secure external authentication over a non–secure protocol such as TCP/IP.

> **Note:** *External authentication* refers to authentication by either a network service or the operating system. *OS authentication* refers to operating system authorization only.

Users request external authentication using the same method they used previous to release 7.2. Users requesting external authentication still use a '/' to indicate the lack of a username and password.

```
SVRMGR> CONNECT /@ny
Connected.
```

If an authentication adapter is available, the server will use it to find the user's network identity. If an authentication adapter is not available, operating system authorization will be performed. If the connection is not secure, the value of REMOTE_OS_AUTHENT will be used to determine if the login may proceed.

> **Note:** It is highly recommended that REMOTE_OS_AUTHENT be set to FALSE, because most protocols are not secure.

Secure Database Links

Proxy authentication (secure database links, that is, secure server-to-server connections) are used when a user who has been externally authenticated by a network service attempts to use an anonymous database link (one without a username or password specified). However, the network service must support some type of credential that can make use of a proxy login. When a connection is requested, the credentials are passed from the local database server to

ORCL01808728

the remote server. The remote server uses the credentials to obtain the network identity of the originating client process.

Network Roles and Privileges

Similar to the secure authentication behavior, if an authentication adapter is linked in, any external roles are retrieved from the network and not the operating system. If an authentication adapter is available (installed and linked into SQL*Net configuration), then it is used to authenticate the user; if not, then the operating system is used. If network roles are supported by the authentication service, they are used.

Configuring external roles varies depending on the authentication service. Instead of using the SID to name network roles (as OS roles does), network roles use the global database name. For example, for DCE authentication, something similar to the following naming convention to configure roles would be used:

```
ORA_<global_database_name>_<role_name>[_[A][D]]
```

> **Note:** Exact syntax to define network roles and privileges may vary depending on the authentication adapter in use.

Remote authentication provides a network version of the OSDBA and OSOPER privileges: SYSDBA and SYSOPER refer to the privileges necessary to perform an internal connection, whether verified by the operating system, password file, or network. A format similar to the following would be used to define network privileges:

```
ORA_<global_database_name>_[DBA|OPER]_SYS
```



SNS

See *Secure Network Services Administrator's Guide* for information on configuring network roles and privileges for specific authentication adapters.

ORCL01808729

ORCL01808730

Exhibit 51

LexisNexis® Total Research System
Switch Client | Preferences | Sign Off | Help
Search | Research Tasks | Get a Document | Shepard's® | Alerts | Total Litigator | Counsel Selector | Dossier | His
Source: News & Business > Combined Sources > Magazine Stories, Combined
Terms: (oracle web server) or (oracle webserver) or (oracle http server) (Edit Search | Suggest Terms for My Search)
Select for FOCUS™ or Delivery

The Rush to the Net;The demand for electronic commerce and advancements in secure transactions are paving the way for full-scale database access via the Internet. Data Based Advisor May, 1996

Copyright 1996 Data Based Solutions, Inc.
Data Based Advisor

May, 1996

SECTION: CLIENT SERVER ADVISOR; Online Development; Pg. 72

LENGTH: 3220 words

HEADLINE: The Rush to the Net;
The demand for electronic commerce and advancements in secure transactions are paving the way for full-scale database access via the Internet.

BYLINE: By Dana Marie Gardner, Associate Editor, and Karen Offermann, Senior Editor; Dana Gardner has a degree in journalism and political science from the University of Richmond, Virginia. She came to Advisor Publications in 1993, where she currently manages the online development section of DATA BASED ADVISOR, and is the managing editor of POWERBUILDER ADVISOR. Karen Offermann spent more than five years with McGraw-Hill's Datapro Research Group as an editor and analyst covering microcomputer operating systems, development tools, and distributed databases. Karen subsequently spent three years as a senior editor at Ziff-Davis' Windows Sources magazine.

BODY:
When the Internet was first conceived, it was a flat-file system used for sharing research materials and documents. There was little interaction; certainly nothing like what you would see today. With the creation of the Wold Wide Web, there came a method to the madness, and the first leg of the Information Superhighway began to take shape.

With recent developments in the IT community, the Internet is sudenly no longer just the amusing diversion it was -- it's become a viable force in business and application development. The Common Gateway Interface (CGI) and developments in security technology have made accessing databases over the Net a reality. Static pages gave way to dynamic ones, giving users up-to-the-minute information on stock prices and weather reports.

Now, corporate users want a piece of the action, and there's a demand for access to corporate data over the Internet. This gives rise to several issues, including how to interface existing databases with the Internet, how to safely make the data available to millions of random users, and how to transmit data from many diverse sources to a sometimes-proprietary client, or Web browser.

Database vendors have watched this evolution carefully, and are rushing to join what has become the next big innovation in computer technology. Not only can the Internet be used

for application deployment, it's widely believed that with the proper security measures, it can be used as the ultimate WAN -- a facilitator for remote computing, synchronization and data sharing between multiple corporate sites, and even as an internal network, or "intranet."

There are already many third-party products on the market or in the works to take advantage of this potential. Recent months have also seen large-scale announcements from some of the largest companies in the industry. Among these are Oracle, Informix, Sybase, Microsoft, and Information Builders, collectively representing the largest chunk of the database server market today.

Oracle

Oracle has been a proactive player on the Net for years. In 1994, the company began distributing its products online. It also has developed an internal network, or Intranet, and an internal version of the World Wide Web. The software giant is among the first to develop a product line for integrating database technology with the Web.

Oracle's march toward integrating database technology with the Internet is focused on using existing technologies to strengthen the Internet's viability as a business solution, and then extending those technologies. The company has based its Internet solution on its flagship Oracle7 database server, which puts it in a position to win all existing Oracle7 users right of the bat.

In September 1995, Oracle released Web Server 1.0, which contsists of Oracle Web Listener (an HTTP server), Oracle Web Agent (a getway to the Oracle7 database), and an Oracle7 Workgroup Server. Then came the announcement of a family of products called WebSystem, at Internet World in Boston this past October. WebSystem includes Oracle's PowerBrowser Web browser, Web Server, and the Web Server Option, which doesn't come with the Oracle Workgroup Server, but does let you connect any Oracle7 Release 7.x database to Web Server. The Web Server Option is the logical choice for customers who already have Oracle7 in-house.

PowerBrowser is Oracle's client, which lets you view information on the Web, and also, to some extent, publish Web pages. "Oracle's goal for PowerBrowser is tobe able to install it on a PC and transform the PC from being a very expensive, dumb client -- when it comes to the Web -- to being something more active," said Randy Hodges, senior product manager of Internet products at Oracle. You can extend PowerBrowser by downloading Network Loadable Objects from the Web, to provide functionality as needed -- for example, a stock ticker for one-time (or extended, if you choose) use.

These Network Loadable Objects (NLOs) will incorporate Sun Microsystems' Java technology, allowing users to download Java-based applets. The Java runtime code, which downloads and excutes programs, is included with PowerBrowser 1.0.

Web Server and PowerBrowser can be used together, or with other vendor's offerings, a feature that Oracle louts as being imperative for survival of, and on, the Net. "The Web has become kind of the universal front end to a lot of different kinds of applications," Hodges said. To provide people with the functionality they need, he said, any Internet solution has to be compatible with othe Internet tools.

Oracle's strategy to keep the WebSystem open includes making Oracle's PowerBrowser interoperable with other browsers, compliance with existing standards -- such as HTTP, CGI, HTMI version 2.0, and some of the more popular features of HTML 3.0, which has not yet been finalized -- and submitting their specifications to Internet standards authorities.

However, despite their effort to remain open, many potential users are hesitant because

WebSystem, at the moment, can't be directly linked to non-Oracle databases. You can connect to other databases using Oracle's gateway products, but according to Hodges, the primary interface to the product is always going to be to Oracle. "You can certainly use the same mechanism with Web Server that Oracle customers have been using for years to get into non-Oracle databases," he said. Users can also use the gateways to connect to ODBC, accessing outside databases that way.

Oracle does have plans to broaden the type of access permitted by WebSystem, Hodges said. The company is planning to release the next version of Web Server, version 1.5, by late March, acording to Hodges. That version will include support for the Secure Sockets Layer (SSL) security protocol, and a way to connect Web Server to other types of applications.

"Generally, we want to open up the back end to a range of services. Oracle led the way in providing integrated database access on the Web, and since then, people have really caught on to the idea. People wanted to incorporate databases to really get the benefits of dynamic database-driven Web pages, and that's what version 1.0 was geared to.

"Now we realize that people want to do more with the Web," he said. "Version 1.5 and 2.0 are really going to focus on providing better performance and better scalability, but also on letting people add in other services to the back end, like server-side Java routines, VB scripts, and CORBA access. People want to get to a number of different types of applications from the Web. That's going to be our openness story over the next few months."

How Web Server works

Applications for **Oracle Web Server** are developed in PL/SQL and stored in the database. Then, when a user comes in from a Web browser and clicks on a URL (address), that request is sent to the Web Listener using HTTP (Hyper Text Transport Protocol). The Web Listener fields the request, looks at the URL, and determines whether it needs to go to the database. If it doesn't, the Web Listener serves up the flat files as with a standard Web site. If the request requires dynamic information, then the Web Listener invokes the Web Agent, which is a CGI application. This invokes the Web Agent process, which calls Oracle stored procedures to connect with the database and execute the PL/SQL. When the PL/SQL runs, it builds the Web page according to the information requested, and sends it back to the browser (figure 1).

In version 2.0, tighter integration between the Web Listener and the Web Agent will let users bypass CGI. "A dedicated pool of database access processes will communicate with the Oracle Web Listener through an IPC. mechanism (an application-level interface), so it's straight process-to-process communication," Hodges said. "Oracle will continue to support CGI for backward compatibility, and so people can write their own programs," Hodges said, "but clearly the way to go is not to use CGI, so you can avoid that performance bottleneck."

Beyond the access

Web Server is being used inside Oracle as an Intranet as well, something that Hodges sees as a major focus of the company within the next year. "The Intranet is our market for 1996, there's no doubt about it. It's the sweet spot," he said. WebSystem, and all its components, can be used as an Intranet solution right out of the box, he said, for such things as information sharing (group-ware), corporate repositories, or a group's home page within a company. "You'll also see the company rolling out a framework of applications geared for such Intranet and Internet use. We'll offer a developrs kit, and also finished applications," Hodges said.

Commerce is another area where Oracle sees an opportunity. By becoming involved with third parties, Oracle is venturing into electronic commerce. The Oracle Web Alliance is a

vehicle for forming these relationships, and involves many companies that cover various facets of Internet technology. With Premenos Corp., which puts out electronic commerce software, Oracle is developing electronic data interchange (EDI) applications over the World Wide Web. AttachMate is also working toward a solution that integrates with WebSystem to allow for business on the Internet. However, Hodges said Oracle is not able to discuss details of either of these ventures at the moment.

Aside from third-party alliances, Oracle is developing a commerce application called Oracle Internet Commerce, which is an Internet sales channel. Dveloped by the Oracle applications division, Oracle Internet Commerce allows users of WebSystem to access internal data from Oracle Applications, which is a suite of tools that includes financial, manufacturing, and distribution modules. The application will let Web pages be updated instantaneously, servicing customers, internal and external employees, and suppliers.

Information Builders

If the issue of openness is to be a deciding factor in the success of Web products, then Information Builders Inc. (IBI) is putting its money in the right place. Known for being able to interact with multiple databases on multiple platforms, IBI's FOCUS product line and EDA product line are at the center of that company's Internet solution.

In addition to the 30 or so platforms that FOCUS and EDA/SQL can develop to, both products can now be used to develop applications to be deployed on the Web. "We think of the Web as a new platform for which FOCUS users can create runtime applications," said John Deutsch, the technical director of EDA marketing. "When you build an application using our tools, you can deploy that application to a Web site."

On the front end, FOCUS and FOCUS Professional will serve as authoring tools. Users will be able to create HTML-ready reports with the products. "With FOCUS, you can press a button to make an application run as a Web HTML application," Deutsch said. You can also make the decision to deploy your application on the Web at the last minute, which means you can create an application and deploy it on any, or all, of the platforms that are supported.

The Publishing end

EDA/SQL fits into the plan as a middleware solution for database publishing, according to Deutsch. The next version, currently in development, is tentatively called EDA Web Connect, and involves all facets of the EDA family of products. Sitting between the Web browser and the back end data sources (figure 2), EDA lends Internet access to non-Web enabled applications.

"We're providing a higher-level, easy-to-use interface through which CGI programs can use SQL to talk to, manipulate, and retrieve information from any of the databases supported by EDA/SQL," Deutsch said. "When you're using, say, a Netscape browser, you'll send a request for information using an URL, and what will be returned to you is a page, or a form, that's been formatted by the HTML language. That form has been created either from a document stored on a Web server, or by a program or executable that lives on the Web Server," he said.

If the request requires data from a source other than the Web server, the executable (a CGI program) is launched and goes through the process of retrieving the necessary information and returning it to the user.

IBI provides three ways of connecting the CGI program with a database. One of these is by using a new high-level command, called the GSQL Entender, Generic SQL (GSQL) is a gateway program that gives you a forms-based interface between Web browsers and SQL

databases. Most authorign tools give developers the option to embed GSQL into their Web forms, Deutsch said. When a CGI program is executed, that program calls the GSQL script. The GSQL Entender then traps the scipt and sends the user's request to the EDA/SQL, which translates it and sends it to the server.

"GSQL, by itself, does nothing," Deutsch said. "It forms a kind of wall socket, which middleware plugs into to translate requests from the CGI program." Recently, GSQL, has become the method of choice for accessing corporate data by most third-party tools, he said. EDA Web Connect also lets you connect to databases using ODBC, or IBI's own SQLAPI, called the EDAAPI.

GSQL is available for use by anyone, so what gives IBI the edge? According to Deutsch, other database vendors can supply like products, but they only work for their own databases or file systems. EDA/SQL provides a SQL interface that can be used for any database on any platform. The access to those remote platforms and databases is provided by the EDA/SQL product. "Effectively, this technology marries applications (CGI programs) with databases and programs being run on traditional corporate machine structure, to hook them up to Web server applications," he said.

For example, if you have an application developed in FoxPro, you can use the gateways and capabilities of EDA/SQL to access the data in that application through a Web page. Users would access your Web site with the browsers of their choice, and send a request to the Web server that has EDA/SQL on it. The request launches a CGI application that resides on the Web Server, which then accesses one of the three APIs for connecting to the database. That API interfaces with EDA/SQL, which translates the request into something your FoxPro database can understand, and the request is processed.

"If a program needs access to traditional corporate databases like Sybase, Oracle, or legacy databases like IMS or VSAM -- in fact, any database supported by EDA/SQL -- the program need only use this product," he said. "It can easily use the SQL language to retrieve or update data on any of the corporation's databases."

The openness angle

Because of IBI's ability to connect to 65 databases on 30 or more platforms, openness is its main edge over competitors, Deutsch said. "The products we have running at customer sites today are similar to what Oracle and Sybase provide," he said, "but when it comes to integrating Web applications with databases, those vendors, for obvious reasons, provide connections with only their own databases. this is what differentiates us."

As with most other database vendors, IBI is supporting the Java language by including a downloadable Java executable in its FOCUS generated HTML applications, according to Deutsch.

IBI is currently using its Internet solution at some client sites for dynamic updates to Web forms. All data is sent in encrypted form, as defined by the Defense Encryption Standards (DES), the Department of Defense's security standards. IBI offers open interfaces for both the client and the servers, so customers can implement their own security, rather than supporting Secure Sockets Layer (SSL) or Secure HTTP (SHTTP).

At the moment, Deutsch said IBI is not delving into the area of electronic commerce, because there hasn't been a need expressed by its users. "Electronic commerce is not initially what our customers have asked us for," he said. They see the use of Information Builders tools as making it easier to build Web applications and integrate them with the databases they currently have and want to reuse as source information as Web server applications."

The GSQL Extender, EDA Web Connect, and the next version of Focus and Focus Professional are expected to be released in the second quarter of this year, Deutsch said.

Sybase web.works

When you consider the issue of DBMS-Internet connectivity, Sybase isn't the first company that comes to mind. Oracle's well-publicized forays into the electronic age have overshadowed most of its competitors, at least in terms of mindset. Yet Sybase has announced a series of tools and services that span the entire range of its product set -- all designed to launch Sybase into the online world: Sybase web.works.

web.works comprises a number of extensions to existing products and the introduction of new products. They include:
* OCX-or class-library-base browser extension to PowerBuilder that allows communication with FTP servers, effectively Internet-enabling existing apps.
* Component Development tools to extend the browser metaphor, providing data access, analysis, presentation, multimedia and business application functionality.
* DataWindow OCX and plug-in.
* Tools to develop executable content, including Java and C++ applets, OLE servers, plug-ins, and OCXs.
* web.sql for connectivity.
* a "to be announced multimedia" authoring tool.

The 1993 Powersoft acquisition gave Sybase a set of respectable client-server developer tools. Since the acquisition, Sybase has moved to integrate Powersoft with the Sybase server line and its other recent acquisitions, including Watcom Software, a SQL engine and C compiler vendor, as well as the development group responsible for the new Sybase Web tools.

The Sybase Internet strategy comprises a series of products and a corollary philosophy. Sybase is one of the first vendors to define server pricing that allows unlimited access to a Sybase DBMS "outside the fire-wall," an attractive pricing model, since it's impossible to sell licenses for Internet DBMSs based on a per-set model. Sybase also supports replication of external DBMSs from within the firewall. Sybase says "As a tool vendor, it's our focus to get people onto the Internet."

Sybase approaches Web content from two perspectives: providing static content in the form of HTML text and graphics, and in the longer view -- creating Web-based executables. The company plans to build extensions to browsers, in the form of Netscape and Sun-based plug-ins and in the form of OLE document servers to handle static data content. Executable content will be supported by building Web-page components with Java applets and Microsfot OLE Sweeper objects.

Standards supported include CGI, Netscape/Sun NSAPI, and Microsoft ISAPI. However, at this time, Sybase has no CGI component analogous to Oracle's Web Agent.

Sybase has also entered into a series of strategic alliances geared toward Internet development. The company is working with Sun Microsystems to define a Java/DBMS access standard. And its recent acquisition of Visual Components launches its new component development model.

GRAPHIC: Figure 1, **ORACLE WEBSERVER** ARCHITECTURE -- Oracle's Internet solution involves using the company's PL/SQL to create on-the-fly Web pages, Figure 2, EDA WEB CONNECT Information Builders' EDA middleware gives any of the 65 supported data sources the ability to be accessed via the Internet;

The tools

Sybase's Internet strategy includes a series of new tools from its two acquisitions: Powersoft and Watcom. Noting "a need for high-level languages and development tools," David Harrah explained that in 1996, Sybase will announce and release a Powersoft-branded product for Web animation. This product will be aimed at users who have a programming skill set approximately equal to that required for developing Excel macros, and won't require a working knowledge of HTML authoring.

Based on a number of open APIs, the Sybase strategy will rely on some OEM-provided components to round out the solution. Harrah indicates that Sybase may OEM a Web server from Netscape.

Sybase's recently announced web.sql provides an interface between the DBMS and HTML., web.sql creates an open, multithreaded link to a DBMS, and maintains a live open connection, providing dynamic content on Web pages. The beta version of web.sql for Solaris is available for download from http://www.media.sybase.com, and is due for release in the first quarter of 1996. Sybase will release the Windows NT, AIX, IRIX, and HPUX versions in the second quarter of 1996.

web.sql developers access data from Sybase System 10 and System 11 databases, and works with leading Web browsers and servers. web.sql lets companies provide dynamic, customized Web pages that reflect individual users' preferences. It also supports selective access to online databases, a feature that could be useful for customer tracking of orders, billing, and so on.

Programmers can use web.sql to link database requests created in SQL or Perl into their web pages, rather than creating and calling a script, and giving applications a potential performance gain. When the page is requested, the query results return to the Web browser as pure HTML text, letting users work with any version of their Web browser.

Because it's capable of maintaining multiple open connections, web.sql can process multiple queries simultaneously. This approach contrasts with Web page programs that must build and break down a database connection for each query.

web.sql works with the Netscape Application Program Interface (NSAPI) or the Common Gateway Interface (CGI), resulting in compatibility with most Web servers and a (slight) performance advantage when used with Netscape Web servers. Its use of CGI-BIN scripts let web.sql work with any standard Web server. Prices range from US $ 1,200 to $ 5,000 per license, depending on the operating system and processor.

Because it incorporates Sybase's Open Client, web.sql exploits Sybase SQL Server, Replication Server, Enterprise CONNECT, and Sybase IQ (figure 3). The combination of web.sql and Enterprise CONNECT lets Sybase web.sql access data in Sybase and non-Sybase databases including DB2, CICS/MVS, Oracle, Informix, and Ingres.

Powersoft Optima++

On February 5th, Powersoft announced its new Optima++ RAD component development tool, designed to let developers gain many of the benefits of C++ component programming without the learning curve. Optima++ is based on the Watcom C++ compiler, and designed for development of Web-based clients and executable applications, as Sun Java applets and Microsoft OLE sweeper controls. Browser extension development is supported using Netscape plug-ins and Microsoft Document Viewers. Common Gateway Interface (CGI), Netscape API (NSAPI), and Microsoft Internet Server (ISAPI) round out the list of industry standards supported.

Optima's drag-and-drop programming metaphor gives developers a path to component assembly, access to Powersoft's client-server features, and a midrange development platform that straddles the skill sets of C++ and PowerBuilder developers. For example, Optima lets you ignore the base C++ class, and simplifies string handling.

The Optima++ tool is geared for an operation in a world of components: those built with Optima, and those that it supports. Sybase will provide some Optima components in the form of .OCXs and Netscape.

The Sybase/Powersoft tools strategy is designed to appeal to programmers at three levels.
* Watcom C++ for system-level programmers (highest performance, greatest programmatic control).
* Optima++ for 3GL component builders and assemblers, and business applications (RAD, lesser learning curve and programmatic control, good-performance and higher than C++ productivity).
* PowerBuilder Enterprise for 4GL and business developers (least programmatic control, highest level of abstraction).

Optima will let programmers create full 32-bit applications, and is focused on Internet, Windows 95, and Windows NT development. If you're a Microsoft Visual C++ programmer, be warned: Optima embraces a different approach from Microsoft's. The MFC classes are not supported, and Sybase describes this product as "component-assembly-based rather than foundation-based."

According to Sybase, the MFC classes aren't supported because they represent a fundamental, non-component way of building an application. While Sybase adds that "if you want, you have the full power of C++," most Windows C++ programmers choose to work with the MFC classes. Any decision to use Optima should be made with this limitation in mind. You can move from Optima to full C++, but only to a proprietary (i.e Watcom, non-MFC) C++ implementation.

Optima will deliver native components linked into .FXE and compiled to the same object driver. Sybase describes this as an efficient way to use OCXs. Theses OCXs are currently targeted to Visual Basic and PowerBuilder. Sybase asserts that support for a dual interface (native binary interface) is expected in many future OCXs. Sybase has briefed a large number of .OCX vendors about its plans and announced development partnerships with several including: Data Technologies, DC Micro Development, FarPoint Technologies, ProtoView Development, and Sheridan Software Systems.

In the Optima programming environment, a reference card will list all the properties you can use with the component (object), and provide a prototype for all components in the OCX.

Optima++ supports all native Windows 95 controls, timers, threads, and standard dialogs in both .DLI and non-DLL versions. OLE automation and components are supported. And for the many of you who know components as VB tools, Powersoft has pledged to make its documented code examples look as much like Visual Basic code as possible.

Basic client-server features supported include data-aware bound controls (query object, transaction object), direct to ODBC drivers, Sybase SQL Anywhere support, and InfoMaker reporting through the DataWindow. The Data Pipeline is also supported.

Optima++ shares a number of technologies with PowerBuilder, including its compiler and code editor: DataWindow, Data Pipeline, and InfoMaker; OLE construction and use metaphor; Object cycle (repository and object management); SQL.Anywhere; and S-Designor.

Optima++ will be available in two versions. The Core Edition for component assembly will support data-aware bound controls, query and transaction objects, ODBC, and single-user SQLAnywhere. The Professional version adds a DataWindows OCX with painters (for ODBC only), the InfoMaker Reporting tool, a three-user SQLAnywhere Server, Data Pipeline ODBC drivers for enterprise DBMSs, Professional C++ with hooks to PVCS, MKS, Object Cycle, and Win32SDK.

Internet access and Sybase security

When you consider Internet access, most IS managers and DBAs express very legitimate concerns about security in this newly connected environment. Tom Parenty, senior product manager, Data and Communications Security at Sybase Corporation, explains Sybase's system security strategy.

The Sybase firewall strategy involves development of Sybase-specific application proxies, letting clients connect to the proxies instead of to the server. Authentication services in the form of public keys, digital signatures, and certificates is expected to appear in Sybase products by the end of 1996.

For Sybase users with integrity concerns, one-way hash functions such as MD5 and the Secure Hash Algorithm provide assurance that a database transaction has not been modified in transit. Sybase uses symmetric algorithms such as DES or RC4 to ensure data integrity.

With increased network connectivity (and its corresponding vulnerability), the DBMS and OS protection mechanisms must be augmented with encryption for some applications. Encryption ensures that only the holder(s) of the key, and not a system administrator, can view sensitive data.

In a time when the use of the Internet as a transport mechanism is increasing exponentially, Sybase views encryption as the only effective way to protect client-server communications. Sybase uses encryption to complement internal DBMS and O/S data security mechanisms.

Informix: Strategic advantages

Informix may just be the sleeping giant of database technology. Renowned for its powerful parallel processing servers. Informix's market share has increased steadily over the past several years, and some analysts claim it's largely at Sybase's expense. Informix products turn up in the most unexpected places, especially when you look hard at the Internet. Parallel processing servers have a great advantage in traffic intensive Web-based applications. Its OnLine Dynamic Server forms the backbone of many of today's most trafficked Internet applications.

Kim Wesselman, Informix senior marketing manager for Internet Solutions explains, "We see a natural progression, from retrieving static data, to retrieving live data, to populating information back to the database." Wesselman adds, "We think it's critical to apply database technology to the Web. Static delivery isn't enough, we need interactivity. By using the Internet and intranets, Informix users can improve their business processes."

Netscape uses the Informix OnLine Dynamic Server in a number of its own turnkey Internet applications. They include: The Netscape Merchant System, Netscape Publishing System, the Netscape Community System, and Netscape Istore.

Unlike its database competitors, Informix is relying heavily on alliances with other companies to enable Internet access. While no one at Informix said "Why reinvent the wheel?", the message was plain. Informix has done very well by relying on its core competencies and avoiding expensive ventures into other, perhaps more glamorous marketplaces.

Informix's strategy is fairly simple: Provide short term-interface kits for existing customers so they can use CGI to provide data access from Web sites. And work with key vendors to integrate "best-of-breed" Web software technologies with Informix's database.

Informix's decision to ally itself with everyone may be the smartest strategy of all. Informix databases work with Netscape Navigator, and both NSCA and Spyglass Mosaic. Informix supports the use of Java for database automation, and of JavaScript to link objects and resources on clients and servers. And, at this writing, Apache, CERN, Netscape, and OpenMarket Web servers also support Informix.

Informix created a Web Partner program in order to merge best-of-class Web-based technologies with the Informix database servers. Concurrent with the unveiling of its Web Partner Program, Informix announced strategic partnerships with several visual Web tool vendors, including Bluestone, Inc., net. Genesis, Inc., Spider Technologies, Inc., and VPE, Inc. These vendors have pledged to provide direct support for Informix's OnLine Dynamic Server or other products in their respective software:
* Bluestone, Inc. (Mountain View, CA; http:www.bluestone.com) will support the OnLine server in its Sapphire/Web visual development tool.
* net.Genesis, Inc. (Cambridge, MA; http://www.netgen.com) will support the OnLine server in its net.Form 2.0 bi-directional dbms access tool.
* Spider Technologies, Inc. (Palo Alto, CA; http://www.w3spider.com) will integrate Spider (its visual Web/database development tool) with the OnLine server.
* VPE, Inc. (Alexandria, VA; http://www.vpe.com) provides WebConnect and Web Builder, each of which supports automated creation of Informix Web applications.

The Illustra marriage

Informix acquired Illustra in December of 1995, with an eye to integrating Illustra's dynamic content management database into the existing core Informix parallel database technology. Successfully adding Illustra's features to Informix's parallel servers will give the company a superb advantage when courting developers who want to deploy data and traffic-intensive Web-based applications (figure 4). But, any acquired technology presents a special challenge when you try to merge it with an existing system.

Illustra's object-relational database handles the melange of text and images that comprise Web pages, and eases the update process. Illustra manages rich media types, including video, combined spatial searches, an extensible multimedia server with cached server options, and enhanced HTML support.

Used by a number of the key vendors in the Internet game, Illustra is the database behind Sun's third-party software catalog, and part of its kernel is built into Navisoft's Web server.

Illustra provides an application template that lets you create a Java class library for your DBMS, so you can access it from the Web. You make the DBMS Java-aware through the class library.

Informix tools

Informix tools

Informix offers free Web interface kits for existing INFORMIX-4GL or INFORMIX-ESQL/C customers who need to develop Web database applications. Each contains source code for writing CGI-Informix database interface programs and example code.

Expected future enhancements to Informix's New Era Toolset will allow you to enable back-

end code access through CGL and permit embedding of a Java application. Details were still sketchy at this writing and subject to change. Kim Wesselman did say the language will be "C++ like." We look forward to hearing more.

If you're planning a Web-based application, and your toolset is still undetermined, take a look at the Informix OnLine Dynamic Server. Once the Illustra technologies get merged into the Informix parallel engine, you'll have a mature Web development system.

Microsoft: Standard operating procedure

Microsoft has never suffered much by arriving late to market with a new technology. In December of 1995, the company announced a complete Internet strategy geared to alleviate concerns about its late entry into the Internet market. All Microsoft products will be enhanced with Internet support, including the integration of the yet-to-be-released Exchange and Internet Information Servers (figure 5). Presumably, the Microsoft Network will become less a proprietary service and more of a Microsoft content site/Internet ramp.

It's odd for Microsoft, standards-bearer extraordinaire, to be in a position where its products seem too proprietary. Microsoft plans to use Connectors to link BackOffice applications with the Internet, and will include an Internet/Database connector with its Internet Information Server. The company will also link the server with the Exchange Server. The proposed Exchange Connector will let users share Exchange public Folders on the Web, and publish Exchange forms to the Web.

On this rock . . .

Two keys to Microsoft's Internet strategy are Gibraltar and Catapult, (secure-access capability). Gibraltar is a World Wide Web, FTP, and Gopher server, due for release by the time you read this. Gibraltar includes a set of filters for logging database traffic requests and an SMTP/Web Exchange connector to connect the server to its upcoming Exchange server, as well as SNMP network management in the server's "shared facilities" module. An SMTP/Web Exchange connector will connect the server to Microsoft's Exchange server.

Microsoft claims that Gibraltar's ISAPI (Internet Service API) will provide better performance for connecting applications to databases than the CGI (Common Gateway Interface) used by most Web servers today. The company offers an ISAPI SDK.

Similar features will also be part of the Internet Business Development Kits, which will include a copy of the Word Internet Assistant that's designed to allow Internet-based documents to be created in Microsoft Word.

Catapult will handle protocol conversion, encryption, and firewall security. It will join disparate networks for Internet access. This protocol-independence will allow users with any IPX at the desktop to connect out using any browser.

Strategic developments

The company's acquisition of privately-held Vermeer was under examination by the U.S. Department of Justice (DOJ) at press time. Vermeer's FrontPage Program lets end users create and manage home pages. While the digital dust settles, in some courtroom, Microsoft plans to integrate Vermeer into a future release of the Microsoft Office Software Suite.

One of Microsoft's key corporate strengths is the ability to attract third-party software developers to support its standards. The VBX and burgeoning OCX markets are prime examples of this philosophy's success. We expect to see a flurry of OCXs designed for use in Internet applications, like Java applets. At least three vendors offer OCXs for embedding

Internet connectivity into the Internet: NetManage Inc., Distinct Corp., and Sax Software Corp. The refocusing of MSN may be the single biggest indicator of Microsoft's move toward the Internet.

ToolTime

Microsoft has announced two programs to encourage use of its development tools for Internet applications. The On-line Internet Software Development Kit lets developers download Microsoft tools for creating Internet applications. ISVs can access authoring tools such as: Microsoft's Blackbird; server software, including the Gibraltar Windows NT server; browsing software, including Microsoft Internet Explorer 2.0; the latest versions of OLE and Internet Services APIs; and libraries of Internet product beta versions.

The Internet Business Development Kit is aimed at helping Microsoft solution providers share information over the Internet. The CD-ROM will include Gibraltar, the Word Internet Assistant (lets users use Microsoft Word as a text editor for the Web), Internet Explorer 2.0, and Blackbird.

A free server pack for the SNA Server 2.11 is supposed to, among other things, simplify access to IBM host data over the Internet. A server-based conversion from FTP to AFTP lets users perform file transfers between IBM hosts and Windows NT servers and workstations without installing the TCP/IP stack on the host. The pack was tested for integration with the upcoming Microsoft Internet Server.

Microsoft's Internet strategy may seem plodding, but odds are that a large number of the Windows installed base will wind up moving ponderously toward Redmond once again. Microsoft's tools will move Windows developers closer to the Web, while leveraging their existing expertise in the slow-to-change corporate PC-based IS environment.

Conclusion

In the fast-paced arena of Internet offerings, it's often the first to market that gets the lion's share. Take, for example, the dominance Netscape has over the browser market, or the way Java has taken the industry by storm. Late-comers not only must get their product out, but also have to go above and beyond the last offering to the extent that they revolutionize the market.

Key offensive tactics in the database/Net wars include supporting Java, creating an alternative to using CGI that alleviates performance bottlenecks, and making products as non-proprietary as possible to maintain the inherently open nature of the Internet. While all the vendors surveyed here are working toward these goals, the strengths of each solution is yet to be seen. (Look for a comparative buyers guide of the products mentioned here in an upcoming issue of *DATA BASED ADVISOR.)*

So far, Oracle has led the database industry to the Net, having released version 1 of their Web Server more than six months ago. However, the race has only just begun. Sybase's impressive stand could pose a threat to this early-established lead, and it's safe to say that you can never count Microsoft out of any marketplace. Information Builders and Informix have also staked a claim on the Internet audience; Informix with several strategic acquisitions and partnerships, and IBI by providing a completely open solution.

With electronic commerce becoming ever more viable through advancements in security, access to corporate data will be an ever-increasing need. Vendors jumping into this market have to do much more than simply leverage their existing user base -- they must aggressively pursue the lion's share of Internet users, and do it quickly.

ORCL01606525

And of course, there's always the possibility that some newcomer will surpass all these, although it's doubtful that any fledgling company could go up against such formidable opponents. Even so, the true fascination of the Internet marketplace is that anything can -- and most likely will -- happen.Figure 1, **ORACLE WEBSERVER** ARCHITECTURE -- Oracle's Internet solution involves using the company's PL/SQL to create on-the-fly Web page; Figure 2, EDA WEB CONNECT -- Information Builders EDA middleware gives any of the 65 supported data sources the ability to be accessed via the Internet; Figure 3, SYBASE WEB WORKS -- Sybase is positioning its new product, web.sql, to work with Enterprise CONNECT for access to non-Sybase databases; Figure 4, INFORMIX STRATEGY -- Informix has chosen to make strategic mergers and partnerships to provide Internet capabilities. The acquisition of Illustra gives new features to Informix's parallel servers; Figure 5, MICROSOFT -- Microsoft is incorporating Internet capabilities into many products, as well as developing new Internet tools. These span from end-user tools to ones for hard-core developers.

LOAD-DATE: May 07, 1996

Source: News & Business > Combined Sources > **Magazine Stories, Combined**
Terms: **(oracle web server) or (oracle webserver) or (oracle http server)** (Edit Search | Suggest Terms for My Search)
View: Full
Date/Time: Tuesday, January 22, 2008 - 2:03 PM EST

Search | Research Tasks | Get a Document | *Shepard's*® | Alerts | Total Litigator | Counsel Selector
History | Delivery Manager | Dossier | Switch Client | Preferences | Sign Off | Help

LexisNexis® About LexisNexis | Terms & Conditions | Contact Us
Copyright © 2008 LexisNexis, a division of Reed Elsevier Inc. All rights reserved.

Exhibit 52

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY

Exhibit 53

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY