

DEVELOPING E-COMMERCE APPLICATIONS USING ORACLE AND JAVA

Megh Thakkar, Quest Software Inc.

INTRODUCTION

The explosion of the Internet has opened up numerous opportunities for the IT professionals. These new ventures have also posed significant challenges that were not encountered before such as:

- Handling of complex data types such as audio, video, HTML and so on
- Protecting data from unauthorized use over the internet
- Performance issues when working with the internet
- Connecting to the database from the web browser
- Providing high availability
- Support for a large number of concurrent users

This paper looks at the challenges faced by companies as they embark on the e-commerce adventure and discusses several techniques provided by Oracle and Java to address these issues.

IMPROVING THE SCALABILITY OF YOUR APPLICATIONS

An e-commerce application potentially has the whole world as its customer base. Special consideration has to be made to ensure that your application will be capable of handling the concurrency and workload. Having said that, we also have to realize that not all these users will be logging on at the same time; however, we have to make a realistic assumption about the amount of concurrency based on the expected workload and the nature of the application. One of the key features of the Java language is the support of multithreading. However, it is still quite difficult to write reliable and scalable multithreaded Java applications. The Oracle8i Aurora Java virtual machine can efficiently and concurrently schedule Java execution for thousands of users. Although Oracle8i supports Java language level threads, all Java code in Oracle8i executes as a call within a session. Therefore, it is not necessary to write server-resident Java code that uses Java's multithreading capabilities to obtain scalability. Instead, the Java programs can use the same scalability architecture used by Oracle internally.

Several solutions are available for increasing the number of concurrent connections:

- Multi-threaded server configuration
- Connection manager
- Connection pooling
- Orastack (for Windows NT only)

INCREASING THE PERFORMANCE

In the Oracle8i Aurora Java Virtual Manager, the memory manager always deals with the allocation and collection of objects within a single session. This enables you to use sophisticated allocation and collection schemes tuned to the types and lifetime of objects, for example:

New objects are allocated in fast and cheap call memory that can be allocated and accessed quickly and efficiently.

- Objects held in Java static variables are migrated to session space.

- Different garbage collection algorithms are applied in the various memory areas, making the overall process very efficient.

REDUCING THE FOOTPRINT OF YOUR APPLICATIONS

For an executing Java program, the footprint is affected by many factors:

- The number of classes and methods and the amount of code they contain.
- Amount of core class libraries used as the program executes.
- The number of objects Aurora allocates, their size, and how many are retained across calls.

In order to achieve scalability and good performance, it is important to have a minimum incremental per-user session footprint. Oracle8I JVM preserves the state of your Java programs between calls by migrating objects reachable by static variables into session space at the end of the call. This indicates that you should use static variables carefully. Improper and unnecessary use of static variables can place a heavy burden on the memory manager and effectively reduce the scalability. Session space usage can be controlled by avoiding the use of static variables for instance-specific data. The EndOfCallRegistry notification provided by Aurora can be used to clear static variables at the end of call.

Several initialization parameters can affect the performance of your Java programs:

- `Shared_pool_size`

The shared pool is consumed transiently when you use `loadjava`. It is also used when call specifications are created and dynamically loaded Java classes are tracked by the system.

- `java_pool_size`

The shared in-memory representation of the Java method and class definitions are stored here. These are shared by all the user sessions. In addition, the Java objects that are migrated to session space at end-of-call are also stored here. The `java_pool_size` must be adjusted based upon the amount of state held in static variables for each session.

- `java_soft_sessionspace_limit` (default 1MB)

Specifies the "soft" limit for the memory usage per session. If this limit is exceeded, a trace file containing a warning is generated.

- `java_max_sessionspace_size` (default 4GB)

Specifies the maximum limit for the memory usage per session. If this limit is exceeded, Aurora kills the session with an out-of-memory failure message. This is a hidden parameter and should not be modified without consulting Oracle support. Improper setting can make your database unstable.

IMPROVING THE SPEED OF EXECUTION

Java executes platform-independent bytecodes on top of a virtual machine. The virtual machine then deals with the specific hardware platform, which makes the execution slightly inefficient. A popular approach to resolving this issue is to use *Just In Time* (JIT) compilers. JIT compilers quickly compile Java bytecodes to native machine code, enabling frequently run Java code to be executed efficiently. Oracle makes use of a *Way Ahead of Time* (WAT) approach. In the WAT approach, Aurora translates Java bytecodes to platform-independent C code. This C code is then translated by a standard C compiler for the target platform. This is suitable for server-side Java applications because they are not updated and deployed frequently. In addition, the WAT approach can be used across all platforms while the JIT approach requires a low-level, processordependent code to be written and maintained for each platform.

Oracle JServer provides all the core Java class libraries and Oracle-provided Java code natively compiled. Java classes exist as shared libraries in the \$ORACLE_HOME/javavm/admin directory. Each shared library corresponds to a Java package and is used by the Aurora virtual machine as needed. In general, the use of natively compiled code instead of bytecode interpretation improves speed (by two to ten times), but it also takes up more memory (two to three times more). The current release of Oracle doesn't provide native compilation for your Java code. In the current JServer release, Java code loaded in the server is interpreted, whereas the core classes such as java.lang.* are fully compiled. As a result, the more Java code used from core classes and the more Oracle-provided class libraries used, the more speed benefit that can be obtained.

IMPROVING SECURITY

A secure environment is very important for continued customer confidence. Lack of security can destroy any organization overnight. Security issues become a bigger challenge in e-commerce applications due to the "access-anywhere" characteristics of such applications. We need to regulate the network traffic that occurs between the company and the internet and vice versa. Attacks from the internet can be thwarted by using a "firewall". In addition to the protection obtained by using firewalls, we also have to secure the database from attacks that can occur within the firewall. Database security is essential so that none of the client applications – whether they are created using Java or any language – can bypass the security and obtain unauthorized access.

When implementing security the principle of least privilege should be followed. This principle basically means that any user should only be given the permissions that are needed for him to perform his duties.

Oracle8i provides various security mechanisms for controlling access to the data. These security features include user logons, passwords, privileges, auditing and fine-grained access control. Quest Software's Quest Central™ can be used to simplify the administration of database security by allowing you to manage users, privileges, roles, profiles and so on.

WORKING WITH FIREWALLS

Most companies that make use of e-commerce would implement an internet "firewall". A firewall is essentially a system that enforces a security policy for data transfer between the company and the internet. In order for a firewall to be effective it is very important that all data transfer between the company and the internet occurs through the firewall. The firewall intercepts, scrutinizes and verifies each and every network packet that passes through it. It is important to understand that the firewall is not a substitute for other security measures. In fact, firewall should be part of a comprehensive security strategy that includes physical protection as well as user awareness of the need to protect their passwords and so on. Firewalls must reflect the level of security used throughout the organization. For example, highly sensitive data should be isolated from the rest of the corporate network and in fact some classified data should not even be connected to the internet.

Generally, a firewall is a combination of hardware and software. Leading firewall vendors make use of one of the following two approaches for implementing firewalls:

- IP filtering

This approach allows access or denies access between networks or specific machines solely based upon information contained in the IP packet headers.

- Application proxy

This approach allows the flow of information through the firewall but not the packets. In other words, direct communication between the inside and outside is not allowed, however, the firewall acts as a data relay between the inside and the outside hosts. You should ideally have some sense of the security policy that will be used during the design phase itself because the type of security used may have an influence on the application design. Specifically, it would be useful to know beforehand whether your application will be transferring data via a firewall or not.

RUNNING APPLICATIONS WITHIN THE FIREWALL

In case your application will run within the perimeter of the firewall, your range of options for communicating between the client and server are extensive – from Remote Method Invocation (RMI) to sockets. Usually the use of RMI also requires the use of sockets access because it makes use of TCP/IP sockets at the network layer.

RUNNING APPLICATIONS THROUGH THE FIREWALL

In case your application will run through the perimeter of the firewall, your range of options for communicating between the client and server are relatively limited – some sort of an HTTP solution. Firewalls typically allow TCP/IP traffic through port 80 – the port traditionally dedicated to HTTP. When implementing a firewall, two diverse philosophies are popular:

- Deny access to anything that is not specifically allowed access

This technique results in high security because it blocks traffic from everything except a few chosen services. Commercial organization would generally implement this type of a firewall.

- Allow access to anything that is not specifically denied access

This technique results in a system that is easy to use but it is not very secure. Academic institutions would generally benefit from this type of a firewall.

When choosing a firewall policy it is important to understand the overall security policy of your organization. Large corporations are usually very strict about the firewall requirements and would not allow you to digress even on an exceptional basis and allow, for example, a small hole in the firewall to allow sockets access. If RMI is used through the firewall then the firewall must allow access to specifically known ports. These ports can't be denied access by the firewall.

UNDERSTANDING THE ORACLE ADVANCED SECURITY OPTION

The Oracle Advanced Security Option provides a range of options for data integrity and security. Some of the functionality is enabled in the Oracle software, however, the majority of the functionality is obtained by integrating Oracle networking with other third-party products. The Advanced Security Option comes in two versions:

- Domestic use version

This version is valid in USA and Canada only and it can use the highest level of encryption currently available

- Export use version

This version is used in other countries and the law requires that the lowest level of encryption be used

The most important functionality provided by the Oracle ASO includes:

- Data encryption and checksumming

Data transmitted between the client and the server can be encrypted so that it is secure during transmission. Further, you can enable the transmission of checksum packets along with data packets so that it is not modified during transmission.

- Authentication and Single Sign-on

The Oracle environment can be integrated with other authentication solutions, for example, Net8 supports adapters for Kerberos, CyberSAFE, SecurID and so on.

- Integration with SSL

Secured Sockets Layer (SSL)- an internet standard can be integrated with Oracle Advanced Security Option

- Support for RADIUS protocol

All devices that comply with the RADIUS standard (Remote Authentication Dial-In User Service) can be used with ASO

- Support for Oracle Wallets

Use of Oracle wallets allows the management of public keys

- Integration with DCE environment

Oracle network and resources can be used with OSF's Distributed Computing Environment (DCE)

ABOUT THE AUTHOR

Megh Thakkar is the Director of Database Technologies at Quest Software. Previously, he worked as a technical specialist at Oracle Corporation. He holds a master's degree in computer science and a bachelor's degree in electronics engineering. Megh also holds several industry vendor certifications, including OCP, MCSE, Novell Certified ECNE, SCO UNIX ACE, and he is a Lotus Notes Certified Consultant. Voted as the best speaker at the NSWOUG'2000 in Sydney and at the VOUG'2001 in Melbourne, Australia; he is a frequent presenter at OracleWorld and various international Oracle User Groups.

Megh is the author of Oracle9i: Instant PL/SQL scripts, e-Business for the Oracle DBA, E-commerce Applications Using Oracle8i and Java from Scratch and Teach Yourself Oracle8i on Windows NT in 24 Hours. He has also co-authored several books, such as Migrating to Oracle8i, Special Edition Using Oracle8/8i, Oracle8 Server Unleashed, C++ Unleashed, COBOL Unleashed, Oracle Certified DBA and, Using Oracle8. Megh is a renowned Oracle specialist who has performed Oracle development, consulting, support, and DBA functions worldwide over the past thirteen years.

REFERENCES:

- Book: "e-Commerce applications using Oracle8i and Java from scratch" by Megh Thakkar (April, 2000) (QUE, ISBN: 0789723387)
- Book: "e-Business for Oracle DBAs" by Megh Thakkar (August, 2001) (SAMS, ISBN: 0672321475)
- Book: "Oracle9i Instant PL/SQL Scripts" by Kevin Loney, Megh Thakkar & Rachel Carmichael (October, 2001) (McGraw-Hill Osborne Media, ISBN: 0072132183)
- Oracle Manuals

MEGHRAJ T. THAKKAR

DIRECTOR OF DATABASE TECHNOLOGIES, QUEST SOFTWARE

MEGH.THAKKAR@OZ.QUEST.COM