




iAnywhere

When is it time to move to
client/server?

Whitepaper

Initial desktop database development

Desktop database applications developed with tools such as CA-Clipper, FoxPro, Visual Basic and Delphi have traditionally been popular because they are relatively inexpensive to implement and deploy. But as more users and Internet and remote access capabilities are added, performance, data integrity and security begin to suffer. When a PC database application is used by only a few users on a LAN, and the database is small, the application usually provides good performance, is easy to maintain, and is relatively stable. Very small user sites often acquire additional tools for use against their database, such as report generators and database browsers. Everything is running smoothly and all database tools are in place. Performance is good and the database is stable.

Increasing users increases the problems

As the number of application users at the site increases, the size of the database doubles, then triples. Transaction loads skyrocket. A remote site is added that requires access to the home site data via a WAN. Suddenly, performance has slowed to a crawl. The network is swamped with traffic and index corruption becomes commonplace. When workstations or the network go down, the database becomes corrupt. With the addition of new users, there is now a need to limit access to key data. Database security becomes a larger issue. Performance is declining. Accessing the database on the LAN is taking longer, and performance of the WAN is even worse—it is unusable. The database is no longer stable or secure. Users have lost their ability to do their work in a timely manner.

Why the performance declined

With desktop database applications, the user interface, business logic, database concurrency control, data movement, data searching, and data manipulation are all performed on the workstation. The file server only acts as a shared hard drive. There is no central point of control to the data. When a workstation needs to search for, read, write, or otherwise access data, it must first attempt to provide database concurrency by locking the data it needs to access. With only a few users on the system, there is little contention for common data. However, as the number of users at a site increases, so does database access contention. The workstation data locking attempts will fail more frequently on initial and subsequent attempts. Not only does the individual application performance suffer but other users' application performance suffers due to the additional network traffic generated by multiple database concurrency lock requests.

The performance problems have become a double-edged sword. The database concurrency locking requires multiple retries and the additional network traffic makes it tougher to make the lock requests.

Why there is more frequent index corruption

The increased frequency of index corruption also has a simple explanation. Every time a table record is to be updated, one or more index files usually need to be changed to complete the update operation. Once an update operation begins, the database will only remain in a stable state if each and every table and index update is completed. Each individual index update operation requires a concurrency lock to be obtained, index page(s) to be read and written, the index header to be read and written, and finally an index "flush to disk" operation to occur. This is for each and every index affected! If the workstation or network goes down at any time during the update operation, indexes will be corrupt, and the database will be left in an unstable state.

Advantage Database Server brings you the benefits of client/server

A move to client/server is in order. What is needed is a product that provides all the security and integrity benefits of client/server and all the performance benefits that should be available in a client/server environment—while allowing developers to keep all their existing applications, their existing hardware, and their years of accumulated expertise. The product that delivers is Advantage Database Server.

Intelligent processing

The Advantage Database Server provides client/server processing by intelligently dividing processing between the client and the server. The user interface and much of the business logic are left to be done by the client workstation. All database concurrency control, data movement, data searching, and data manipulation are performed by the Advantage Database Server on the file server. The Advantage Database Server becomes a central point of control for database access. Unlike non-client/server desktop applications, Advantage moves the shared data processing operations to the file server where it makes sense—where the data is stored. Only the results of the database operation are returned to the client. The superior architecture of Advantage leads to far better performance than non-client/server systems.

Eliminate concurrency contention problems

Because the Advantage Database Server is the central point of control to the database, data concurrency contention issues are minimal. The intelligent lock management system of the Advantage Database Server allows multiple users to search through and read from shared files concurrently. This leads to drastically improved performance in multi-user systems. By performing all database concurrency control and data access on the server, network traffic is reduced significantly which also leads to drastically improved performance in multi-user environments. Reduced network traffic not only helps LAN performance, but it makes WAN database access feasible whereas before it was out of the question.

Eliminate index corruption

Advantage Database Server eliminates index corruption. Table and index updates are performed on the server. A table or index is never passed to the client to be updated. When a table record and the associated indexes are updated by Advantage, the update operations do not occur until all necessary information has reached the Advantage Database Server on the file server. Even if a workstation or the network goes down, no index corruption will ever occur and the database integrity will always remain intact.

Transaction processing

The Advantage Database Server also adds a complete Transaction Processing System (i.e., Begin, Commit, and Rollback Transaction) to your application. This allows application developers to define business transactions which usually involve multiple updates to one or more tables in the database. If one or more of these updates in the defined business transaction do not occur, the database is left in an unstable state. With the Advantage Transaction Processing System, all of the updates necessary to complete the business transaction would be placed between the Begin and Commit Transaction statements. If the workstation or network were to go down before the Commit Transaction was reached, all updates would be automatically rolled back. If the server were to go down before the Commit Transaction was reached, all updates would be automatically rolled back once the server was brought back up, and the Advantage Database re-started. In either case, the database would be left as if the transaction did not even begin.

Referential integrity and database constraints

Advantage Database Server provides Referential Integrity and Database Constraints to verify the validity of data in your database and maintain relationships between records in your database. Specifically, Advantage has support for Unique and Primary Keys, Referential Integrity, Field-level constraints, and Record-level constraints.

Security

The Advantage Database Server can be used to provide database security by allowing only Advantage-enabled applications to access the database. The system administrator can take away all desired network access rights to the directories that contain the database tables and index files, so that non-Advantage applications cannot gain access to the database. Unauthorized access to the database is eliminated. Advantage Database Server provides further database security functionality by allowing creation of user accounts to control the access to tables, columns, views and stored procedures.

Use what you already have, add to what you already know

The Advantage Database Server brings client/server functionality to existing database applications. No application re-write is necessary. Advantage provides plug-and-play client interfaces for the most popular development platforms, including Delphi, C++Builder, Kylix, Visual Studio .NET, Visual Basic, Access, Java, Perl, PHP, CA-Clipper and CA-Visual Objects. Since you do not have to change applications or tools, and no user training is necessary, existing system expertise can continue to be used. The Advantage Database Server does not require new dedicated server hardware. It runs on existing NetWare, Windows NT/2000/2003, or Linux servers. Advantage Database Server installs in minutes and is easy to maintain. It does not require a database administrator for maintenance like other database systems. Advantage also allows developers to convert their existing applications one at a time to use the Advantage Database Server, eliminating lengthy down time during integration. If your site has multi-user performance problems, suffers from data corruption, and lacks database security, it is time to look for a client/server solution. Advantage Database Server is simply the best and most cost effective client/server solution available. Advantage provides the performance, database integrity, and database security expected in client/server systems without requiring application re-write, new hardware or loss of existing expertise.



ANYWHERE SOLUTIONS, INC.
WORLDWIDE HEADQUARTERS
ONE SYBASE DR IVE
DUBLIN, CA 94568-7902
U.S.A.

FOR GENERAL INFORMATION:
AdvantageInfo@iAnywhere.com
NORTH AMERICA
T 1-800-801-2069
1-519-883-6898

FOR SPECIFIC REGIONAL
PRODUCT INFORMATION:
UNITED STATES
(800)-235-7576 ext. 5030
(208)-322-7575 ext. 5030
BENELUX
+31(0)73 6235359
GERMANY
+49 (0)7032/ 798-200
UNITED KINGDOM
+44 (0)117 333 9000

AdvantageDatabase.com

ANYWHERE SOLUTIONS IS A SUBSIDIARY OF SYBASE, INC. COPYRIGHT © 2006 ANYWHERE SOLUTIONS, INC. ALL RIGHTS RESERVED. ANYWHERE, SYBASE, AND THE SYBASE LOGO ARE TRADEMARKS OF SYBASE, INC. OR ITS SUBSIDIARIES. ALL OTHER TRADEMARKS ARE PROPERTIES OF THEIR RESPECTIVE OWNERS. ® INDICATES REGISTRATION IN THE UNITED STATES OF AMERICA.
L02808

iAnywhere
A SYBASE COMPANY