Paper# 3014

HP IMAGE/SQL Today

Kumar KN
Database Lab
Commercial Systems Division - India
Hewlett-Packard Company
(408) 447-5489

This paper represents the collective work and efforts of several individuals from the Database Lab and Marketing organizations of the Commercial Systems Division of Hewlett-Packard Company.

****

1 INTRODUCTION

For those of you who are new to IMAGE, IMAGE is a database management system (DBMS) from Hewlett-Packard available on the HP3000 system. Since its introduction in the mid 1970s, it has been winning awards for its function, design, and most of all, its performance and reliability. Thousands of applications, from email to finance to production control, all operate using IMAGE as their database management system. The number of users executing programs against data stored in IMAGE databases is, no doubt, very large.

Without argument, IMAGE gives you high reliability and high performance. After more than twenty years of improvements and refinements, you, as customers and users, would expect nothing less. With so much going for this product, with a dependability that few other database products offer, as a storage management tool that is a de-facto standard on HP3000 computer systems, what could possibly be done to improve it? Read on for the answer.

2 HP IMAGE/SQL

Hewlett-Packard is proud to talk about HP IMAGE/SQL, a relational access to IMAGE data using industry-standard Structured Query Language (SQL). This method of access includes full read and write capability using ANSI standard functionality. This new access method makes a myriad of application development and decision support tools available to IMAGE users. With this, you can also deploy IMAGE in client-server environments with remarkable ease.

An exciting part of this whole story is that SQL access to IMAGE and the access that you enjoy today are totally compatible. Complete coexistence! No conversions, no recompilation, no changes. The only change will be the way you look at your data once you use the application development and decision support tools now available with IMAGE. And these are the same tools that can be used with all the big name relational databases, including ALLBASE/SQL. IMAGE has now moved into the Open Systems arena.

Those of you familiar with IMAGE will undoubtedly ask the question, "but what will this do to performance?" Of course, every database design and application will have it's own performance characteristics, but, as a general rule, IMAGE/SQL will perform at about seventy to ninety percent of native IMAGE access. The term 'native' here refers to intrinsic

level access using COBOL, PASCAL, et cetera. So, you will have the best of both worlds, fast native access and relational access.

With IMAGE/SQL you will have everything that you have today, wrapped in a new relational package, with an enormous selection of new tools, without any hassles. Let's take a look at the issues HP has tried to address with IMAGE/SQL.

3 TOP ISSUES

There are a multitude of issues which HP hopes to resolve with IMAGE/SQL, and these will be explained shortly. Unfortunately, some new issues will be created as well, and these will also be mentioned. These issues have been grouped according to the kind of organizations that will be impacted, and they are VARs (Value Added Resellers) and ISVs (Independent Software Vendors), MIS or IT (Information Technology) Departments, and end-users.

Most VARs and ISVs, whose products include applications and tools, have been limited to non-relational application development tools when dealing with IMAGE. There are also very few PC-based GUI (Graphical User Interface) products which offer transparent interaction between PC clients and HP 3000 servers using IMAGE. As mentioned before, thousands of applications have been developed by numerous companies all using IMAGE as the storage management system. Now, with IMAGE/SQL, all new SQL-based application development and decision support tools, most using some type of GUI, some operable in a client-server environment, can be used by VARs and ISVs to improve and enhance their products. This new transformation of IMAGE data into relational data will revitalize the VAR/ISV product environment.

Turning to MIS and IT departments, the primary issues are:

- reducing the application backlog
- providing flexible information access
- recruiting SQL-trained personnel

Each of these will be elaborated below.

Ideally, IMAGE/SQL should reduce the application development backlog for any MIS/IT organization. This reduction would be accomplished by using any of the easy-to-use application development and decision support tools now available. The opposite effect, that of increasing the demand for new applications, will also probably become true. As the end-users see the amazing results from new implementations using these tools, all kinds of previously hidden application requests will surface. Once the power of IMAGE/SQL is known, MIS/IT departments will be flooded with new requests for information.

Does this mean that these departments should avoid using IMAGE/SQL just to avoid this rampage? Absolutely not! The gains which a company can realize from improved information review and analysis, as the result of improved data availability, can be significant. Users will be able to do their own queries and reporting, MIS need to only supply the access. Some of the other issues in this area will further explain these statements.
The second issue in this area pertains to flexible information access. What this refers to is the powerful, command-driven language which characterizes SQL. Complex and

sophisticated queries can be executed to supply end-users with their specific information requirements.

SQL uses sentence constructs of SELECT and WHERE clauses, and the relational JOIN clause, to define its syntax. This level of knowledge is not necessarily required to use the various tools, however, more sophisticated applications might require SQL expertise. Finding the right people to develop these queries and deliver these applications to the end-users will be critical. This brings us to the last issue for the MIS/IT department -- finding qualified personnel.

Relational concepts in information management have become popular in the last several years. These concepts have become the de-facto standard for data storage systems at educational institutions all over the world. As such, many computer professionals have had either some or extensive exposure to relational concepts.

SQL has become the relational language standard and is taught widely. Finding qualified personnel to program using SQL today is easier than finding experienced IMAGE programmers. So, even though your information investment is in IMAGE, the investment you make in accessing that information can be in relational technology.

The issues of the end-user are improved productivity and improved decision support. Both of these issues rest firmly on the same foundation; information that the end-user requires to be more productive, and upon which decisions are made, must be located where the user can exploit and present in a meaningful form. This can be achieved with the use of PC-based client/server tools featuring graphical user interfaces (GUI). These are the very same tools which can be used by the MIS/IT department for development, and VARs and ISVs to enhance their offerings.

4 DATABASE TOOLS

The new tools for relational databases have been mentioned in the previous paragraphs. Before getting to specific products, let's first describe them in general, and then divide them into categories to better understand what they can do and offer.

All of these tools are PC-based and function primarily in a client/server environment operating with Microsoft Windows. With exception, these tools use ODBC, which is based on a standard introduced by Microsoft, for connecting to the server.

For those of you not familiar with what an API is, some explanation is in order. API stands for Application Programming Interface and refers to the component of the client/server model which performs the interactive linking between client and server over the network for the purpose of data exchange. This can be thought of as application level handshaking, analogous to RS232C which is an electrical signal handshaking in data communications.

Our first category of tools is Decision Support. As the name suggests, the purpose of these tools is to support business decisions, and this is accomplished through information analysis, reporting and graphical representation. Typically, this type of tool is oriented towards end-users doing financial and managerial analysis where numerical quantification and graphical representation of data is useful. These tools also tend to offer formatting of data for subsequent importing to spreadsheet products for further

manipulation and review. End-user knowledge of SQL is not a prerequisite for using these tools.

Some examples of Decision Support Tools are Impromptu by Cognos, and HP's Information Access/iX. Although this last tool is not new, the relational access now delivered through IMAGE/SQL will dramatically reduce the overhead which will improve administration of these systems and make this solution easier to setup and maintain.

Another category of Decision Support Tools is spreadsheet use such as MS Excel and Lotus 1-2-3. These products support the ODBC interface and thus can directly access data stored in IMAGE/SQL databases.

The last category of tools for use with IMAGE/SQL is Application Development Tools. Falling also into the category of Fourth Generation Languages (4GL), these tools offer PC Windows programming capability with Graphical User Interfaces which allow programmers fast development of critical end-user applications. In many cases, knowledge of SQL is not necessary.

Where Decision Support and EIS tools tend to be read-intensive or read-only, application development tools are intended to create interactive applications for online transaction processing. An example in this category is PowerBuilder by Powersoft.

All the aforementioned categories make up the new tools available for use with IMAGE/SQL. These tools also operate with HP's ALLBASE/SQL and other major independent relational databases.

5 CURRENT 4GL TOOLS

In addition to all the PC client/server tools mentioned above, a rich assortment of direct-access 4GL tools exist for accessing HP's relational databases. All the tools in this section work with both ALLBASE/SQL and IMAGE/SQL. The important point about these tools is that they execute on the host, not the client/server cooperative execution of the previous set of tools.

The products in this section have, for the most part, been around for some time. Details about features and benefits will not be presented here. The main point here is that several vendors offer SQL-based tools for accessing relational databases.

Some of these 4GL tools offer native (intrinsic) access to IMAGE. However, the relational access which they offer to ALLBASE/SQL now extends to IMAGE/SQL too. Below is a list of products and the companies which offer them:

ALLBASE Toolset	Hewlett-Packard
Transact	Hewlett-Packard
Powerhouse	Cognos
Focus	Information Builders
JAM	JYACC
Speedware	Speedware
Uniface	Uniface

As you can see, these tools, along with those indicated for PC client/server, make an impressive arsenal in your development efforts.

Both relational databases from HP, ALLBASE/SQL and IMAGE/SQL, can reside on your system, and all these tools can access either or both.

6 HP ALLBASE/SQL

HP offers two database management systems on the HP 3000 platform. Each has its place with any given application, and many applications could use either. The majority of database usage on the HP 3000 is currently IMAGE, but there are certain applications where ALLBASE/SQL is the better choice.

ALLBASE/SQL is a full-featured relational database management system (RDBMS) with functional compliance of ANSI standards. HP sees four specific areas where ALLBASE is the preferred DBMS:

Mainframe class computing
Distributed applications
High-volume online transaction processing
Object-oriented applications

HP has positioned ALLBASE as an alternative to mainframe-class database management systems. This works in conjunction with HP's mainframe downsizing strategy which offers high-end HP 3000 systems as Corporate Business Systems. Providing online backup and restructuring, and supporting very large file sizes, ALLBASE is the DBMS of choice in this area.

Distributed information and the applications which support them are also addressed within the ANSI standards for relational database systems. Here again, ALLBASE has been specifically featured. Using two-phase commit protocols, distributed transactions can be ensured of atomicity.

ALLBASE also supports Encina technology from TRANSARC Corporation. This technology provides a standardized method of transparent distributed transaction processing in a multi-platform environment. These distributed application feature makes ALLBASE the clear choice compared to IMAGE/SQL.

Where high-volume online transaction processing (OLTP) is the objective, and relational concepts are required, ALLBASE is again the best choice. This pertains mostly to new application development. Since ALLBASE has been designed from the ground up as a high-end relational database, applications requiring high-volume SQL OLTP will reap the benefit.

Object-oriented computing is being seen more and more as a viable solution in many applications. Most assuredly, any application which has multimedia requirements can benefit from object technology. Here ALLBASE/SQL has the advantage. Object storage, especially that of video, photograph, audio, et cetera, requires data structures foreign to IMAGE/SQL. ALLBASE/SQL already provides the necessary storage with Binary Large Objects (BLOBs).

6.1 COEXISTENCE

The HP3000 now offers a complete range of data management choices. Native IMAGE still provides the highest performance in the industry for mission- critical OLTP business applications. IMAGE/SQL provides data access through the multitude of 4GL and PC client/server toolsets at a very small performance premium. Client/server computing is clearly the emerging trend in information processing, and IMAGE/SQL is now an important element of it, ensuring the protection of your information investment.

ALLBASE/SQL provides the highest SQL performance of any relational database in the industry and provides support for distributed database and distributed transaction processing. With Corporate Business Systems, ALLBASE/SQL can handle the requirements of nearly any enterprise at a fraction of the cost of traditional mainframe solutions.

Together, these database management systems make a powerful team. And they work together. Concurrent access is possible by linking the two together within a single environment. Your applications, whether host-based or client/server, can simultaneously access ALLBASE/SQL and IMAGE/SQL information. That is coexistence at its best.

But there is more.

6.2 THE RELATIONAL PICTURE

HP offers you two relational databases, IMAGE/SQL and ALLBASE/SQL. But the story gets better. Also available on the HP3000 is a link to IMAGE/SQL through the Oracle Transparent Gateway for IMAGE/SQL. This creates a SQL shell over every database management system available on the HP 3000. And with this shell comes all the 4GL and client/server tools which have been mentioned throughout this paper. This is a very powerful offering, indeed.

7 COMPONENTS OF IMAGE/SQL

IMAGE/SQL includes TurboIMAGE/XL, a restricted version of ALLBASE/SQL, and a database administration tool called IMAGESQL that links them together. Each of these components is described below.

TURBOIMAGE/XL

> A set of programs and procedures used to create and maintain IMAGE databases.

ALLBASE/SQL

> A set of programs and procedures used to create, maintain, and access relational database environments (called DBEnvironments or DBEs). This includes various components which are required to provide relational access to IMAGE databases. ALLBASE/PC-API is software that provides access from a Microsoft Windows based PC client to ALLBASE/SQL and IMAGE/SQL databases. ALLBASE/NET provides transparent, remote access to IMAGE/SQL and ALLBASE/SQL databases.

IMAGESQL

A database administrator's tool used to manage IMAGE databases in a relational environment. This tool registers information about the IMAGE database in the ALLBASE/SQL DBEnvironment, thus paving the way for front-end tools to use this information to access IMAGE databases.

INFORMATION ACCESS SERVER

This is the HP3000 server component of HP's Information Access/iX and is bundled with IMAGE/SQL.

8 HOW IMAGE/SQL WORKS

Accessing IMAGE data relationally is simple and consists of two steps.

Step 1: Attach the IMAGE database to the DBEnvironment
Step 2: Access IMAGE data relationally

Each of these steps will be briefly described next.

Step 1: Attach the IMAGE database to the DBE

The attach process examines the IMAGE database and places equivalent SQL structure information in the DBE. This step uses the IMAGESQL utility.

```
:IMAGESQL
> SET TURBODB  CUSTDB
> SET SQLDBE  SQLDBE
If the specified DBE does not exist, you will be prompted for its creation.

> ATTACH
```

The above example attaches an IMAGE database named CUSTDB to a DBE named SQLDBE.

Step 2: Access IMAGE data relationally

There are several methods to access data relationally. The following example uses ISQL.

```
:ISQL
ISQL=>  CONNECT  TO  'SQLDBE';
ISQL=>  SELECT * FROM CUSTDB.CUSTOMERS;
```

The result of this query would be a tabulated list of all entries in the CUSTOMERS data set. You can also display information about the database itself.

```
ISQL => SELECT NAME, OWNER FROM SYSTEM.TABLE;
```

This would return all table names associated with the SQLDBE DBEnvironment.

Updating IMAGE databases is also simple as shown below. This ISQL example updates a column (item) called PRODUCT_NAME.

```
:ISQL
ISQL=> CONNECT TO 'SQLDBE';
ISQL=> UPDATE CUSTDB.ORDERS
        SET PRODUCT_NAME = 'IMAGE/SQL' WHERE PRODUCT_NAME = 'IMAGE';
```

The UPDATE statement finds all entries where PRODUCT_NAME is 'IMAGE', then changes that value to 'IMAGE/SQL'.

As the above steps demonstrate, accessing IMAGE data relationally is very simple and straight-forward. Once the database is attached to the DBE, nothing else is required in the regular use of the system. There are additional administrative tasks related to security and data type mapping for the purposes of fine tuning but these are not complex either.

Next, we will describe some of the enhancements made to IMAGE/SQL which serve to make coexistence much better than what it seems.

9 ENHANCED IMAGE/SQL

The following enhancements to IMAGE/SQL have significantly improved performance and made the task of using IMAGE/SQL much simpler with the new functionality.

Indexed access in IMAGE/SQL

> The ALLBASE/SQL optimizer is now aware of IMAGE key items, search items, and third-party indexes. As a result, data is retrieved more efficiently when the mapped column represents one of the above items or indexes. With this, you can leverage your third-party indexing functionality from IMAGE/SQL. All information about the indexes is registered in the ALLBASE/SQL DBE when the IMAGE database is ATTACHed to the DBE.

IMAGE K8 data types

> The UPDATE TYPE command of IMAGE/SQL is now enhanced to allow update of your K8 data type to one of the SQL data types DATE, TIME, DATETIME, or INTERVAL. By doing this, dates, times, and intervals can be entered through the SQL interface and stored directly in the IMAGE database.

Multi-connect feature

> The Multi-connect feature in IMAGE/SQL enables you to simultaneously connect to more than one database environment.

Predicate level locking

Predicate level locking, instead of set level locking, is used when possible to provide a high level of concurrency.

10 CONCLUSION

HP IMAGE/SQL is one of the most significant enhancements to the IMAGE DBMS, as you may have deduced by now. The opportunity which this product provides to commercial and company developers is enormous considering the selection of tools now available for use with IMAGE. IMAGE/SQL also solidifies its coexistence with HP's other relational database, ALLBASE/SQL, with standard access to both and new improvements are being made to better this coexistence.

The simplicity with which relational access occurs with IMAGE/SQL makes this perhaps the most painless 'conversion' in the history of the computer industry. The term 'coexistence' applies equally well to the dual accesses to IMAGE, native and SQL, as it does to the two relational databases, IMAGE/SQL and ALLBASE/SQL.

With IMAGE/SQL, HP has once again protected your investment in the HP3000 and offers you the means to access it using industry standard SQL!