



ACCESS ISAM FILE DATABASES

INTRODUCTION

Open Database Connectivity (ODBC) compliant access to data has become a primary requirement by end users for vendors supplying applications containing data. It is becoming even more crucial as corporate organizations accelerate the deployment of client/server computing. Widely used tools such as word processors, spreadsheets and database application development tools are ODBC enabled and they provide means of building powerful end-user applications quickly through ODBC access.

Many vertical applications use ISAM (Indexed Sequential Access Management) type of files for the storage and management of data. These applications require all data access to be performed using customized programs and tend to exist on SCO UNIX, VAX/VMS, Alpha/ OpenVMS or Windows NT type of platforms. End users are interested in accessing this information from the desktop for reporting and application. More recently, the ability to make this information available on the web provides a new paradigm for data access and distribution.

In order to make this information easily accessible and usable, one needs a software infrastructure that:

- ▶ Supports access to information stored in proprietary databases from standard desktop tools like Microsoft Access, PowerBuilder, Visual Basic, Crystal Reports and many other tools. Today this means ODBC and JDBC compatible access.
- ▶ Allows you to quickly develop client applications for the PC environment using any of the hundreds of desktop tools.
- ▶ Can be developed in a very short time without becoming an ODBC, JDBC, OLE DB or SQL expert.

Once a user has the ability to use ODBC to access your ISAM database system, the possibilities are endless.

ACCESSING ISAM FILES – A SAMPLE CASE AND REQUIREMENTS

For this sample case, assume we have a Widget Management System (WMS) that runs on a UNIX platform. Users access this system from terminals attached to the UNIX box. Currently all interactions with the WMS are done by logging into the UNIX box from terminals and running the WMS programs. The only way to transfer data between the WMS and a PC is by exporting a report into a file and then transferring it to a PC. Today, users of the WMS want to get this data directly into Microsoft Excel and Crystal Reports. The WMS stores data using a CISAM compatible database.

HIGHLIGHTS:

- ▶ Quickly implement an interface to Indexed Sequential Access Management (ISAM) Files
- ▶ Execute SQL statements
- ▶ Formulate queries in the language of your tool
- ▶ Implement business rules

We want the following features:

1. ODBC/OLE DB access - To provide read/write access to all or some of the data files.
2. Fast data access - Optimized access to the files which may contain millions of records.
3. Optimized join processing - Handle joins between many files that may contain millions of records.
4. Flexibility in schema management - Allow the use of either the existing data dictionary feature in the database or easily set one up using a supplied schema manager.
5. Support for future desktop data access standards - Support JDBC, OLE DB and other technologies that may come up in the future without changing the server.
6. Support all my platforms - Need server support for Alpha/OpenVMS, VAX/VMS, Intel/NT, Alpha/NT, RS-6000/AIX, Intel/SCO UNIX, HP-9000/HP-UX, SPARC/Solaris, NCR-3300/AT&T UNIX and IBM MVS. Want to implement code on one platform and easily use on all others.
7. Ease of distribution - The ability to easily distribute the server and client components.

HOW TO QUICKLY IMPLEMENT AN INTERFACE TO YOUR ISAM FILES

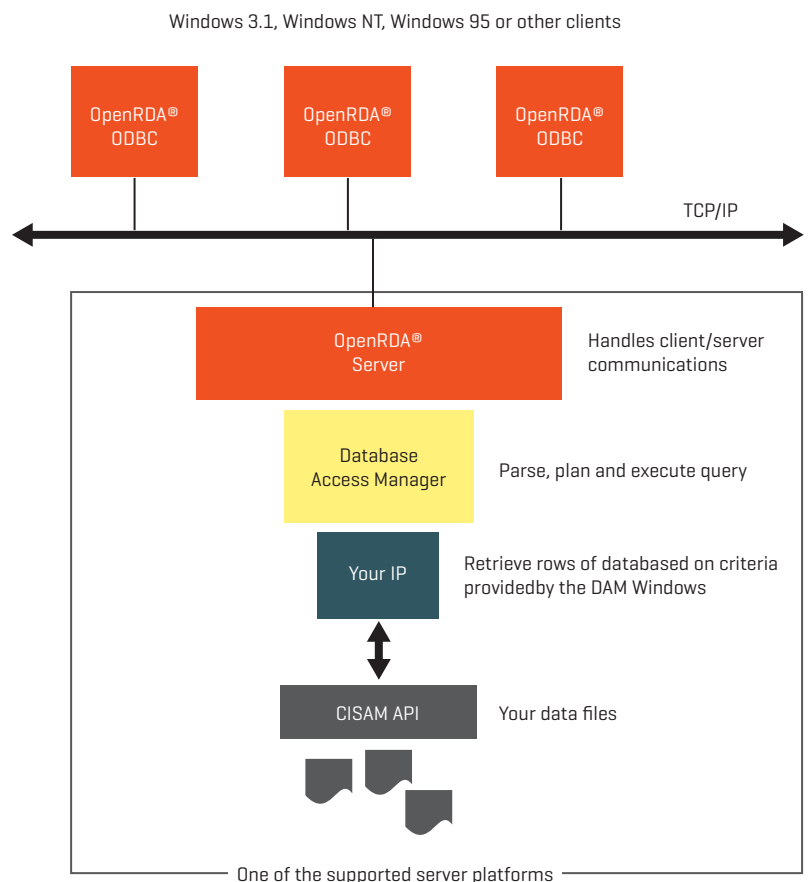
The Progress® DataDirect® OpenAccess™ SDK product can help you create a client/server or local ODBC/OLE DB interface to your database within weeks. We provide 90% of the code in binary form. You just have to implement the functions to read and write rows of data from the files in your database—no need to learn ODBC, OLE DB or SQL. Figure 1 shows the proposed architecture using the OpenAccess SDK product. The only portion you develop is the box labeled “Your IP”. This code is responsible for accepting a data access request along with the optimization information from the Database Access Manager [DAM] and using this information to build a set of rows.

EXECUTING A SQL STATEMENT AGAINST THE ISAM DATABASE

Let’s walk through the sequence of steps we need to perform in order to execute a SQL statement against the ISAM database. The following steps describe the operation the IP performs in the course of executing a query:

1. Receive connection from the PC client with a user name and password - Use this information to provide access verification to your database system. This information can be used to restrict different classes of users to certain table access. If security does not exist, this information does not have to be used. The IP receives this request when the client issues an ODBC connect call.

Figure 1: OpenAccess SDK based architecture



2. Receive an execute request - The IP is called to execute a query when the client issues an ODBC SQLExecute call. At this time, the IP can call the functions in the DAM to find out what table is being accessed and what conditions are specified in the query. For example, take the query given below.

```
select PART,MFR,DESCRIP from PARTS where PARTNO = '1FM1' (Q1)
```

To execute this query, the IP will first determine the table being accessed and the requested operation. In this case the operation is SELECT and the table is PARTS. Assuming the IP has been set up to use index information, the IP first checks to see if any condition on an indexed column is present. If so, the IP asks the DAM to report restrictions on that column [PART in this example]. In this case the DAM will report [=,'1FM1'] as a condition on column PART. The IP can find out if additional columns in the query are indexed and then use this information to get restrictions on those columns. The IP then uses all this information to call functions to retrieve the data from the ISAM files. In our example, we use the function WMS_positionInFile[fileName = 'parts.dbf',indexColumn='PART', value='1FM1'] to position the file pointer to a record where the PART=1FM1. We then call WMS_readRecord[recBuf] to read the record at the current position into the recBuf buffer. For each row retrieved from the database file(s), the IP builds a row and passes it to the DAM for further evaluation and for placing in the result set. To build a row, the IP breaks apart the data record in the rcBuf buffer into individual field values and maps them to the columns in the row.

There is no need to have one table to one file mapping. The IP is effectively exposing views to the client and each of these views can be made up of multiple files. Exposing data from multiple files into one table is sometimes better than requiring the end user to always issue joins. Joins are less efficient to process than single table queries and they complicate the writing of queries.

Receive a disconnect - The IP frees any resources allocated on behalf of the client.

As seen by the above steps, the OpenAccess SDK makes it very simple for the IP to access data in an optimized way. The above steps can be implemented in less than 50 lines of code.

FEATURES SUITED FOR ISAM DATABASES

In this section we highlight some of the features of the OpenAccess SDK that make it ideal for accessing large or small databases maintained in ISAM type files.

Specifically, the OpenAccess SDK provides:

- ▶ Access to full details of the query
- ▶ Optimized join processing
- ▶ Schema management
- ▶ Transaction control
- ▶ Ability to support stored procedures
- ▶ Ability to implement business rules

ACCESS TO FULL DETAILS OF THE QUERY

The DAM allows the IP to find out as little or as much as it wants about a query. This helps the IP use information contained in the query to limit the number of rows it reads. Information about the query can be obtained on individual columns or on the entire WHERE clause. You can retrieve the entire WHERE clause as a set of expressions on the referenced columns. This information can be used to reformulate the query into your own syntax if your database supports a query language. You can also retrieve the full expression information to implement powerful optimization schemes.

OPTIMIZED JOIN PROCESSING

Joins are very expensive unless handled properly. The DAM handles all joins. The IP is only responsible for accessing data from a single table at a time. For join processing, the IP is called to process a SELECT for each table. The simplest way to implement a join is for the IP to execute the query on one of the tables, execute the query on the next table, and then perform a Cartesian product between the two tables. A Cartesian product of two result sets of size M and N is M*N. This is a very large number as M and N grow and this happens very often when a single piece of information is spread amongst many tables. The DAM optimizes this by first building a result set for the first table and then going through each row in this set and passing the required information to the next query as restrictions. This way the off diagonal elements are not even created. The IP is called to process SELECT on the second table as many times as there are rows in the first result set.

SCHEMA MANAGEMENT

The OpenAccess SDK allows you to provide the schema information (data dictionary) for your database tables by either populating the Schema Database provided as part of the DAM or by implementing schema function in the IP.

TRANSACTION MANAGEMENT

The OpenAccess SDK allows you to perform full transaction control by passing the start transaction and end transaction events to the IP. It is the responsibility of each IP and the database to maintain logs required to implement commit and rollback. Also, we support the SELECT ... FOR UPDATE syntax to allow you to implement row level locking of records.

STORED PROCEDURES

OpenAccess SDK allows client applications to invoke stored procedures within the IP. This allows the IP to implement many functions on the server as part of the IP code. A stored procedure is implemented as part of the IP code or within your database system and exposed by a name and the required arguments.

BUSINESS RULES

OpenAccess SDK allows the IP to easily implement business logic. Since an IP is exposing a view, it has complete control over how selects, inserts, updates and deletes are processed on each of the tables. The IP can take advantage of this to validate the queries and to guarantee that business rules are enforced. You also can easily check the validity of the input data and control the allowed queries.

YOUR DEVELOPMENT EFFORT

- ▶ Define the schema for the data to be exposed [2 man days]
- ▶ Implement the IP code to access the tables defined in item 1 – start with our sample [2 man days]
- ▶ Test with ODBC applications [2 man days]
- ▶ Optimize [3 man days]
- ▶ Package up for distribution [2 man days]

Expected time of completion: **12 man days**

Expected time for working prototype: **2 days**

DESIGNING AND CODING THE ISAM IP

NOTE: The full source code for the ISAM IP discussed here is available on our web site at the same location as this use case document. Please contact DataDirect sales or support if you did not receive it. Follow the instructions in Installing the Sample ISAM IP in this document once you have obtained the required files.

Implementing an IP requires two major tasks: 1) schema management 2) data access management. We will first discuss how the schema needs to be set up to allow reading/writing from ISAM files, and then we will discuss the implementation of the IP code that will process the ISAM files to expose the data as per schema definition.

SCHEMA DEFINITION

One of the first things to consider in designing an IP is how the data dictionary will be implemented. OpenAccess requires the IP developer to define the schema for the databases using one of two methods:

1. Populate the built-in schema database.
2. Implement schema function in the IP code to handle requests for looking up table, column, index, and other information.

The first method is what our documentation refers to as static schema and is suitable for applications where the schema is not changing because the end users cannot add or modify table definitions. The second method is suited for databases where new tables are being created or existing ones are being modified by the existing application.

In either case, the schema database has to provide the required information to map data from the records in the files to the columns of a table defined for that file. In our example the PARTS table is exposing data in a file parts.dat with an index file parts.idx. The records in the parts.dat file are divided into fields with each field being of a certain size and containing certain type of data. The position of each field and the type of data contained in that field is needed as part of the schema information to allow the IP code to map records from the file into a row required by the SQL engine.

In the schema, for each table we can use the oa_userdata field to supply the associate file name. In our case, for the table PARTS, we will set oa_userdata to parts.dat. This information will be used by our isopen() function to open the file. For each column definition we can use the oa_userdata to store information about the native data type and position within the record. For example, the definition of column PARTNO would have: 'oa_userdata = 1;4' to indicate that column PARTNO starts at position one in the record and is of type integer.

IP CODE

The code that implements the IP needs to be generic in that it uses information stored in the schema database to process the specified queries. This way new tables can be handled without modifying the code.

NOTE: All functions starting with isam_ are part of the ISAM IP and all functions and typedefs starting with DAM or dam are part of the OpenAccess SDK libraries.

The isam_ip_execute() function for this IP would look like the code shown in Listing 1. It is called to process the rows for each table involved in the query execution. First we call dam_describeTable to find out what table is being queried and to obtain the file name that has the data for that table.

Next we find out what all columns are referenced in the query. For each of these columns we obtain the mapping into the file record structure. This is done by the isam_getColumnsInfo function. This function makes use of DAM functions that allow you to walk through either all columns in the table or the columns that are referenced in the query.

Next we call `isopen()` to open the file associated with the current table being accessed. What we do next depends on if we need to do a full table scan or not. To do this we call `dam_getOptimalIndexAndConditions` to check for any conditions on columns that are marked as indexed. We want to pick up cases such as 'where PARTNO=xyz'. In this case the above function call will return a condition list `[=,xyz]` and we can use this information to jump to the location in the `parts.dat` file where this matches. Depending on your database design, we can also handle cases such as 'where PARTNO > xyz' by jumping to the record where index value is xyz and then reading all rows whose key values are > xyz.

Assuming the `dam_getOptimalIndexAndConditions` returns a = condition, we can use `isstart()` function with the key value to find the first record that matches the key. If the index is marked unique then we are done; otherwise we have to keep reading next records until no records matching that key are available. The logic here is to walk through each of the index values and build rows for it. The function `isam_buildKeyInfo` is called to find out information about the columns that are part of the index. The function `isam_buildKeyRec` is called to build a key record for the selected condition list. This function uses the information about the column to file record mapping to build up a key record that can be passed to the `isstart()` ISAM function. For a query of the form:

```
select * from PARTS where PARTNO=1 or PARTNO=2
```

We would have two condition lists. First we would process all rows with PARTNO=1 and then all records with PARTNO=2. Once we have moved to the record that matches the key record, we read the record using `isread()` and then call `isam_buildRow` to convert the record read from the file into columns of a row that can be passed to the OpenAccess SQL engine. Next, we call `dam_isTargetRow` to see if the current row matches the full where clause. If it does then we add it to the result set. Now if the index is not defined to be unique, then we need to scan rest of the rows after the current row until we find a non-matching row. This logic can be enhanced to handle BETWEEN, >=, <= and LIKE operators on indexed columns.

If `dam_getOptimalIndexAndConditions` returns no conditions, then we have to do a full table scan. This happens if the query contains no conditions on columns that are part of an index. In this case we call `isstart(ISFIRST)` to position to read the first record in the file and continue building and processing the rows until we are at the end.

We have discussed handling select here but implementing update and delete is just a matter of adding extra code right after when `dam_isTargetRow` is called. If true is returned then you either delete or update the current record. Insert is handled simply by getting the rows to insert from the OpenAccess SQL engine and adding them to your file. Again you would make use of the column to record field mapping to build the record to write into the file.

ISAM INTERFACE

The sample IP explained in this application note makes use of standard ISAM API calls as listed in Table 1. In the program listing these functions are highlighted with shading.

ISAM Function	Description
<code>isopen</code>	Open the specified file
<code>isclose</code>	Close the specified file
<code>isstart</code>	Position the record pointer to the specified record. A key can be specified. By default the first key for the file is assumed. Key value can be specified.
<code>isread</code>	Read the specified record.

Table 1: ISAM API Used in This Example

INSTALLING THE SAMPLE ISAM IP

Please note that the ISAM IP will not compile and link unless you provide the required ISWRAP.H files and the required library that implements the ISAM API.

1. Obtain source code from the DataDirect Knowledgebase
2. Create a directory isam under the damp directory
3. Create a directory src under the isam directory and uncompress the source files in it.
4. Now follow the instruction in the IP Examples document you downloaded or received as part of the SDK. In the instructions use ISAM for the IP type and substitute all occurrences of {example} with isam. The ISAM IP uses static schema.

WHAT'S NEXT?

This use case provides overview and details of using the OpenAccess SDK to implement an ODBC or OLE DB interface to your flavor of ISAM database. You can start with our sample ISAM IP and add:

1. Support for update/delete/insert
2. Cursor based select processing
3. Stored procedures for commonly performed business logic
4. Other data types

PROGRESS SOFTWARE

Progress Software Corporation [NASDAQ: PRGS] is a global software company that simplifies the development, deployment and management of business applications on-premise or in the cloud, on any platform or device, to any data source, with enhanced performance, minimal IT complexity and low total cost of ownership.

WORLDWIDE HEADQUARTERS

Progress Software Corporation, 14 Oak Park, Bedford, MA 01730 USA Tel: +1 781 280-4000 Fax: +1 781 280-4095 On the Web at: www.progress.com

Find us on  [facebook.com/progresssw](https://www.facebook.com/progresssw)  twitter.com/progresssw  [youtube.com/progresssw](https://www.youtube.com/progresssw)

For regional international office locations and contact information, please go to www.progress.com/worldwide

Progress, DataDirect, DataDirect Connect, OpenAccess, and SequeLink are trademarks or registered trademarks of Progress Software Corporation or one of its affiliates or subsidiaries in the U.S. and other countries. Any other marks contained herein may be trademarks of their respective owners. Specifications subject to change without notice.

© 2008, 2014 Progress Software Corporation. All rights reserved.

Rev. 9/14

www.progress.com

