

by Lisa Clark
Senior Programmer
BellHawk Systems Corporation

Introduction

This article presents the pros and cons of various databases commonly used on the factory floor. It covers databases such as Access, SQL Server and Oracle. Its intent is to enable people implementing factory-floor information systems to make an informed choice.



Access 2000

Microsoft Access 2000 is a powerful relational database application with which a desktop user can efficiently create and manipulate database systems. It contains an application development environment for VB for Applications programmers that includes forms technology, reports, and database administration. There is also a user interface common to both Access and the other Office applications. The combination of ease of use and power in Access makes it a top choice among developers.

Access is often used as a front-end application that can process data stored on a back end. In this scenario, the Access application runs on the client machine, accessing data stored on a database server that's running software such as Microsoft SQL Server. Access does an excellent job acting as the client-side, front-end software in this scenario. But what is your best choice for the back end server? This article will consider four options for the back end: Access, MSDE, SQL Server and Oracle. We will also briefly discuss the hardware required for implementing a database solution.

Access as the Back End

It's possible to use Access as both the front-end and the back end, however, the Access database isn't acting as a true backend in that it's not doing any processing. Many people mistakenly believe that an Access MDB database file stored on a file server acts as a database server, but that's not really the case. Imagine that you have a table stored in an Access database on a file server that contains 400,000 customer records. The user wants to see a list of all the customers in Massachusetts who placed an order within the last 6 months. With the data stored in an Access MDB file format, all 400,00 records are sent over the network to the workstation and the query which selects only the desired records is performed on the workstation. This results in significant network traffic.

On the other hand, assume that these records are stored on a database server such as Microsoft SQL Server. The user runs the same query. In this case, the server performs the query processing and only sends the names of the Massachusetts customers who meet the criteria over the network to the client workstation. This situation is referred to as client/server processing.

The following issues should be considered when deciding whether to upsize to a true client/server configuration:

- As the volume of data within an Access database increases, you'll probably notice degradation in performance. Many people say that 100 MB is the magic number for the maximum size of an Access database, whereas many back end database servers can handle databases containing multiple gigabytes of data.

- Just as a large volume of data can be a problem, so can a larger number of concurrent users. In fact, more than 5 users concurrently accessing an Access database can really degrade performance.
- Certain applications demand better performance than other applications. Imagine 100 users simultaneously taking phone orders; it would not be appropriate for the users to wait 15 seconds between entering each item that's ordered. On the other hand, asking a user to wait 60 seconds to process a management report the user runs once each month isn't a lot to ask. Not only does the client/server architecture lead to better performance, but most back end database servers can utilize multithreaded operating systems with multiple processors. Access can't.
- As a file server within an organization experiences increasing demands, the Access application might simply exacerbate an already growing problem. By moving the application data to a database server, the reduced demands on the network overall might provide all users on the network with better performance, even those not using the Access application.

Using MSDE as the Back End

With the introduction of Access 2000, users and developers have a choice of data engines. They can continue with an improved version of the default Jet data engine and MDB files, or they can use MSDE. The MSDE is a free, de-featured version of Microsoft SQL Server that ships with Microsoft Office 2000 and Microsoft Access 2000. This tool is part of Microsoft's future strategy of moving Access off the Jet engine and entirely onto SQL Server architecture. Microsoft offers a user-friendly interface to the MSDE in the form of an Access data project (ADP). This new file type is tailored to work with MSDE and other SQL Server style databases. Access projects seamlessly integrate remote SQL Server data sources with Access forms, reports, modules, and data access pages that are stored locally in the ADP file. Access projects offer faster, less resource-intensive data access to SQL Server data sources. Users don't have to link to a remote data source or cache data locally by linking an ODBC data source, as was the case with prior Access versions.

MSDE incorporates technology from SQL Server 7.0. SQL Server delivers a single code base, which scales from a PC running Windows 95 to multiprocessor clusters running Windows NT Server, Enterprise Edition, offering 100 percent application compatibility. The older Jet database engine, which was shipped with previous versions of Access does not have this type of scalability.

The big advantage that MSDE enjoys over Jet is that it is a client-server data engine; Jet is a file-server data engine. Because Jet is a file-server system, the query processing must happen on the client. This involves moving a lot of data over the network for large databases. MSDE runs that same query on the server. This puts a larger load on the server, but can reduce network traffic substantially, especially if the users are selecting a small subset of the data. Another advantage that MSDE has over Jet is that MSDE is a process that logs transactions. If anything should go wrong while writing to the database, such as disk error, network failure, or power failure, MSDE can recover because it logs transactions. After the system comes back up, MSDE will revert back to the last consistent state. This gives MSDE greater reliability than Jet. If the system were to go down with Jet, the database could be corrupt and you might need to revert back to your last backup copy.

To summarize, both Jet and MSDE are best when you have a small number of users.

Jet Access is more appropriate when:

- You want the highest compatibility with Access 97 or earlier.
- You have very low resources, such as memory or disk.
- Ease of use is a primary concern.

MSDE is more appropriate when:

- You expect a future need for greater scalability.

- You require easy merge replication with the central server.
- You need the best security.
- You need great reliability, such as transaction logging.
- Your system is online 24 hours a day, 7 days a week.
- You need stored procedures and triggers.

Using SQL Server or Oracle as the Back End

How do you know when it's time to invest in a full featured database server like SQL Server or Oracle? These systems not only cost more in the first place but also require a greater investment in the systems professionals necessary to run them.

Both Jet and MSDE can handle up to 2 GigaBytes of data, however, the full SQL Server version 7.0 enables individual files to grow to 32 TeraBytes and permits up to 32,767 files per database. MSDE supports up to two symmetric processors on Microsoft Windows NT systems while SQL Server 7 can recognize up to 32 multiple processors. SQL Server 7.0 can also handle a very large number of simultaneous users while Jet and MSDE are optimized for individual or small workgroup solutions with a maximum of 5 concurrent users.

The backup and recovery options offered with an Access database stored on a file server simply don't rival the options for backup and recovery on a database server. A decent database server has very powerful uninterruptible power supplies. Many have hot swapping disk drives with disk mirroring, disk duplexing, or disk striping with parity (redundant array of independent (or inexpensive) disks (RAID) Level 5 protection). Disk mirroring and duplexing mean that data can be written to multiple drives at one time, providing instantaneous backups.

Furthermore, some database server tape backup software enables backups to be completed while users are accessing the system. Many offer automatic transaction logging. All these mean less chance of data loss or downtime. With certain applications, this type of backup and recovery is overkill. With other applications, it's imperative.

Access offers what used to be considered the best security for a desktop database, but it can't compare with that provided by most database servers. Database server security often works in conjunction with the network operating system. This is the case, for example, with Microsoft SQL Server and Windows NT Server. Remember that no matter how much security you place on an Access database, this doesn't prevent a user from deleting the entire MDB file from the network disk. It's very easy to offer protection from this potential problem on a database server. Furthermore, many back end database server products offer field-level security not offered within an Access MDB file. Finally, many back ends offer integrated security with one logon for both the network and the database.

Choosing between SQL Server and Oracle will partly depend on what you already have. If you are a Windows only shop, SQL Server is fine. The Oracle solution would be better with mixed platforms. In general, Oracle products are designed for very professional development efforts by top-notch programmers and project leaders. The learning period is fairly long and the solution is pricey, but you will ultimately have greater scalability and greater reliability. Microsoft offers a solution that's aimed at rapid development and low-cost implementation. The tools are cheaper and the servers to run it on are cheaper. The key problem with SQL Server is that you'll be tied to Microsoft operating systems and Intel hardware.

Recommended Equipment Configurations

For the database server, any Windows server operating system (such as Windows NT 4) running an MSDE or SQL Server 7 or later database is an excellent choice. The minimum recommended configuration for an NT Server box is a 300 MHz or faster Pentium with 256 MB of RAM and a 10 GB

hard drive. But memory is cheap these days and the more the better. In a high usage environment you might want to go up to 512 MB or possibly 1 GB.

It is recommended that external SCSI mirrored hard drives be used so that the server can continue operation in the event of disk failure. Some organizations prefer to use RAID array disk drives to further enhance systems reliability. There are all sorts of RAID hard drive configurations that will give you reliability and performance. Of course they all cost varying amounts of money and that's where you make the trade off. The advantages in reliability are that mirrored hard-drives allow you to continue running even though a drive fails. The more sophisticated setups allow you to hot swap a defective drive... pull out the bad drive, put in a new one, and the system takes care of everything. The users never notice a thing and there is zero downtime. For performance (and somewhat for reliability), there is disk striping that essentially will take an array of hard drives and write in parallel to all of them. This allows you X times the thruput where X is the number of drives (minus some overhead for managing all of that).

The server should be equipped with floppy and CD ROM disk drives, as well as a keyboard and a mouse. A small 15-inch monitor is adequate for systems maintenance purposes. An uninterruptible power supply is essential for power conditioning and controlled shutdown during power outages. It is important that the server be equipped with a mechanism for frequent backups of your database. This should either be a local tape backup unit on the server or a network-based backup from some other computer. The server should be connected to the shop floor LAN by means of a 10Mbits/sec, and preferably 100 Mbits/sec, connection.

Your other decision involves the client computers on the factory floor that will be running the front-end application program. You will need to check the system requirements for your software but generally for most Access front end programs the computers should be 100MHz or faster Pentium processors running Windows 98 release 2, Windows 2000 or Windows NT operating system. They should have at least 64 MB of RAM, at least 500 MB of Disk storage, and a 10 MB/sec or faster connection to the factory floor LAN. In general, 15-inch monitors are adequate for shop floor use. A minimum resolution of 800x600 pixels is necessary for shop floor use but 1024x768 is beneficial if managers are to use the shop floor computers to view the supervisor's screens.

Conclusions

If you are expecting less than 5 concurrent users, the database of choice is MSDE, which ships with Access 2000. For more users, the database of choice is SQL Server. If your staff's database maintenance expertise is entirely focused on Oracle, then Oracle becomes the database of choice.

Copyright © 2001 BellHawk Systems Corporation, an operating unit of Applied Real-Time Intelligent Systems Corporation. BellHawk ® and WebHawk ® are registered trade marks of BellHawk Systems Corporation.

Go to www.bellhawk.com for more information.