



ADO AND ADO.NET DATABASE ACCESS TECHNOLOGY COMPARISONS

Vaikunth Pai

Department of Information Technology, Srinivas Institute of
Management Studies, Pandeshwar, Mangalore, Karnataka

Abstract:

The ADO.NET is a brand new data access technique of the Microsoft. It is the bridge of the .NET application and database communication, which use the XML to exchange data and add many new program interfaces to access database. This paper aims to shed some insights for Microsoft web developers to understand the two technologies in depth. An overview of ADO and ADO.NET from a relational perspective is presented, with particular emphasis on the ADO Recordset and the ADO.NET Dataset objects. A comparative analysis is performed providing an assessment of strengths and limitations of each technology. Conclusions drawn show that ADO.NET presents some significant advantages over its predecessor such as a superior disconnected model with greater control over the update process, improved management for multi-table processing, as well as improved scalability and object model extensibility.

Index Terms: ADO, ADO.NET, Record Set, Data Set, Relational Database, Scalability & Extensibility

1. Introduction:

Access, transport, storage and manipulation of data are fundamental and underlying processes at the heart of every application/system. ADO and ADO.NET support these same processes, albeit through different supporting run-time platforms, infra-structures, classes, levels of functionality and part of different Microsoft strategic paradigms.

ADO is a data access technology that provides a high-level easy-to-use interface to access relational data as well as other types of data. It is a programming and data-access model that represents a framework in the form of an object oriented programming interface. ADO is based on Microsoft's proprietary COM architecture, a building block that evolved application systems from 2-tier, to N-tier, to distributed computing frameworks and then to distributed applications that incorporated the Internet.

ADO.NET is part of the .NET initiative introduced in June 2000 as a new strategy for the development, deployment and execution of highly distributed applications referred to as Web Services. A networking infrastructure is provided that is layered on the Internet or Intranet. This infrastructure is leveraged against technologies that are both ubiquitous and nonproprietary, and built upon industry set standards and protocols. The .NET initiative provides a completely new .NET platform – a new .NET Framework development environment and a new Common Language Runtime run-time environment. ADO.NET, by default, derives many of the benefits of being part of the .NET environment such as a foundation based on the acceptance of an object-oriented approach, coupled with a new self-describing component model – the .NET assembly. Some important benefits include a rich class library that is extensible and organised by namespace, a Common Type System, managed execution of applications, language interoperability, structured exception handling and simple deployment.

The next two sections provide an overview of ADO and ADO.NET object models. These are then used as a basis for delivering comparisons and conclusions in Section 4 and 5, respectively.

2. ADO:

ADO is a class set providing access to relational data. It provides transparency using a common interface to access various data sources using a core set of objects such as Connection, Command and Recordset. There is also Parameter, Property, Field and Error, a set of subsidiary objects that are accessed via Parameters, Properties, Fields and Errors collections respectively. All of these objects and collections have a Properties collection that holds property objects. ADO objects feature properties, methods and events to be able to manipulate the object and its contents.

Using set of objects, ADO provides flexibility in choice as to how to access and update a data source from a client application. For example, Recordset properties offer a matrix of choices such as for Cursor Location (Client, Server), Lock Type, (ad Lock Read Only, ad Lock Batch Optimistic, ad Lock Optimistic, ad Lock Read Only) and Cursor Type (Forward, Static, Key-Set, Dynamic). The setting of these properties needs to include consideration of connected or disconnected states and cache size.

The inherent ADO flexibility creates overlapping in the use of some objects to perform the same action, or different actions performed by the same object e.g. sharing common objects for different access/update approaches. There is some potential to not select the optimal combination for the desired purpose. For example, ADO allows changes to be applied directly to a data source using the following objects and methods:

- Connection object Execute method
- Command object Execute method
- Recordset object Open method

Alternatively, updates may be performed using an in-memory connected or disconnected Recordset object using a server-side or client-side cursor, with read-only (ad Lock Read Only), pessimistic (ad Lock Pessimistic), batch optimistic (ad Lock Batch Optimistic) or optimistic (ad Lock Optimistic) lock types. Cursor type is always ad Open Static if cursor location is client-side. A Recordset with Batch optimistic locking allows changes to be applied without any data source interaction, even if the Recordset holds an active connection.

An optimistic locking Recordset allows Recordset rows, individually or as batch, to be updated without holding a lock on the corresponding record at the data source. Success or failure in propagating a modified data row to the data source depends entirely on whether another user has since modified that same row. If a Recordset uses optimistic locking (ad Lock Batch Optimistic or ad Lock Optimistic), data access and updating is stateless. Choosing ad Lock Batch Optimistic allows Recordset updating without needing to use system resources to maintain a connection. Alternatively ADO Recordset data access and updating may be stateful when connected with a pessimistic lock.

A Recordset with a client-side or server-side ad Lock Optimistic cursor allows a row that is pointed to by the cursor to be updated or cancelled before proceeding to another row. If the cursor is set to ad Lock Batch Optimistic, changes may be made wholesale within the Recordset without data source interaction. These changes may then be reconnected to, and updated against, the data source or cleared from the Recordset.

3. ADO.NET:

Whereas ADO provides a multi-faceted approach to data handling, the ADO.NET model provides a philosophical change in approach where its set of objects

are functionally optimised for specific tasks. ADO.NET data access is based upon the Dataset and data provider that provides specific access to a data source.

The Dataset is akin to a mini-relational database, but stored in memory. It is composed of a core set of objects that include its namesake, as well as the Data Table, Data Column, Data Row, Data View, Data Relation and Constraint objects. The Dataset is equated with the disconnected aspect of ADO.NET. The Data Set is a standalone object existing independently of any data source.

The data provider provides the means to communicate with any data source. A .NET data provider provides a set of objects that enable the connection to, data retrieval from, data update to, and disconnection from, a data source. Each data provider is for a specific data source. Each data provider provides its own set of classes such as for SQL Server .NET, OLE DB .NET or Oracle .NET. As an example SQL Server .NET uses Sql Connection, Sql Command, Sql Parameter, Sql Data Reader, Sql Data Adapter and Sql Transaction objects (located in System. Data Sql Client namespace) and OLE DB .NET data providers use Ole Db Connection, Ole Db Command, Ole Db Parameter, Ole Db Data Reader, Ole Db Data Adapter and Ole Db Transaction.

The ADO.NET model introduces a level of abstraction where there is a separation between data and the process by which data is retrieved and updated. The centerpiece of this model is the data adapter. Each data adapter has its own specific data provider. It is the bridge between the Data Set and the data source. It is the coordination layer between the in-memory Data Set and the permanent sources of data access through the data provider.

The data adapter uses the data reader of the same data provider to populate a Dataset using the Fill method. The data adapter uses four instances of the Command object: Select Command, Insert Command, Update Command and Delete Command, to propagate, via the Update method, the changes of the same type in a Dataset to a data source e.g. an inserted row in the Dataset is used with the Insert Command. It does not automatically include this logic. The filling of a single Data Set may be performed using multiple data adapters. The population of a Dataset is not limited only to data providers but may be sourced from user input, programmatically or XML input. The data adapter does not interact directly with the connection object but indirectly through its set of four command objects. The data adapter always leaves the state of a connection as it was, that is opened or closed, after executing. This is a default process that supports the disconnected paradigm of the ADO.NET model. The Connection object provides access to a data source either to retrieve data, or to provide the pipeline for applying data changes to a data source.

.NET provides data providers. These are also referred to as managed providers such as OLE DB .NET and SQL Server .NET. The OLE DB .NET managed provider provides access to various data sources requiring COM interoperability services of the .NET Framework to access OLE DB functionality. A managed provider such as SQL Server .NET provides a more efficient, direct and streamlined access to the API of the database, in comparison to the OLE DB provider. For example, the SQL Server .NET managed provider uses Sql Data Reader. Which uses its own Tabular Data Stream protocol. This protocol is managed by the CLR to access data where SQL Server exposes its data types to conform to the .NET Framework types as well as being able to expose its data in native format. That is, when retrieving results from a managed Data Reader such as Sql Data Reader, columns are retrieved in their native data type form without requiring expensive conversions.

4. ADO And ADO.NET Comparisons:

Both ADO and ADO.NET provide similarities with their command objects allowing retrieval from, and update of, a data source. However as noted above, ADO provides multiple ways to update a data source directly, whereas using ADO.NET, the only way to update a data source directly is by using the command object of a data provider. ADO.NET provides alignment of Command object methods to the type of result returned. The methods offered include Execute Reader, Execute Non Query, Execute Scalar and Execute Xml Reader. This compares to the ADO Command object where only one method, the Execute method, is used, and where the Command Type property determines the type of result expected. The Command Type property may be assigned an enumerated value representing values such as a stored procedure, table name or Sql query.

ADO.NET command types may be text i.e. an SQL string, a parameterized stored procedure with input or output parameters or a table name(s). Additionally ADO provides access to data from persisted files or data in semi-structured form.

Furthermore, the Connection object is only associated with the ADO.NET Command object whereas with ADO the Connection objects is associated with the Active Connection property that is available to both the Recordset and Command objects.

In ADO.NET, the Command object Execute Reader method may be used to return a result into the data providers data reader object. In comparison, and as described earlier, ADO may return a Recordset result set from a database using the Recordset, Command or Connection objects.

The ADO.NET data reader object provides a fast-forward read-only data stream that may be read to access row data consecutively using the Read method. It cannot be instantiated – it is always provided through the command object interface. However it is limiting if random access is required. It may instead be stored in an array or populated into a Dataset using the data adapter. Data Reader objects are essentially the equivalent of the forward-only read-only Recordset cursor of ADO.

.NET managed data providers use separate libraries and classes for different data stores. ADO uses the same Connection, Command and Recordset set of objects for different data sources. This is in comparison to ADO.NET, which provides a different group of objects for each data provider. This means that using more than one data provider requires more than a simple connection string substitution the exception being OLE DB .NET data provider.

ADO represents a single unified object model attempting to be all encompassing for all data sources - ADO.NET does not.

5. Recordset and Dataset Comparisons:

Compared to the ADO.NET Dataset object, an intimate understanding of the Recordset is required to obtain maximum benefits when choosing between a servers or client located cursor, the degree of locking required and the level of visibility desired to changes applied by other users to the data source. The overlapping in usage of the Recordset does not provide a clear delineation between using disconnected and connected models. This provides the potential to not attain optimal benefits when designing access to, and update of, a data source.

Although the Recordset and Dataset provide similarities, there are also some compelling differences. The Data Set may be identified with a Recordset that has a client-side Cursor Location, a Cursor Type that is open static and a Lock Type that is

optimistic. However, the Data Set extends the capabilities of the Recordset for managing application data.

Both an ADO Recordset and an ADO.NET Dataset may be created manually, implicitly or explicitly, from the schematics of a database. Both in-memory models enable the storage of tabular/relational type data and which may be populated in various ways. ADO.NET provides the Data View object which is based on a Data Table object. Using the Data View object, data may be modified, sorted and filtered with the capability to clear changes or to apply changes permanently to a data source e.g. database. There is no equivalent of this object in ADO. Both facilitate the creation of a smaller Recordset or Dataset that is an extraction containing only changes. This provides a lower transmission load when transmitting changes across the network to update a data source.

In comparison to the Dataset, the Recordset is limited to a single tabular result. The Dataset resembles a relational database system i.e. it may have multiple tables, relationships between tables through column associations, primary and foreign key referential integrity, update and delete cascading control. Of course a Data Table may hold a result set derived through a JOIN SELECT statement. This may be appropriate at times when a non-hierarchical view of data is sufficient, e.g. through a grid control. A Recordset that contains a SELECT JOIN result set is difficult to synchronise.

The Dataset is completely data source agnostic, whereas the Recordset is not. Once a Dataset is populated e.g. from a data provider, it has no knowledge of where the data came from - it has no direct interface with the Connection object. This is in contrast to the Recordset object that has a 'connection' to a data source through the Active Connection property. When using the Data Set, the connection is only open when retrieving data from a source, or propagating changes to a data source.

A Dataset may be typed or untyped. A typed Dataset provides several advantages over the Recordset such as providing compile-time checking, Intellisense, as well as code that provides easier programming – it is more readable and concise.

6. Table Relationships:

The Recordset may be populated with a result set from an SQL data source representing a single table (or subset of). If only a single table is retrieved, the Recordset and Dataset have much in common.

When there is a requirement to access data from two or more tables, differences arise between the use of a Dataset and a Recordset. If a Recordset receives a result set from a join of two or more tables, two points of interest arise:

- The presence of duplicated data
- Ascertaining the source of the data is not straight forward

Although reading and updating is straightforward, extra programming effort is required to control the insertion and deletion of Recordset rows against root tables. Base and unique tables need to be established, otherwise undesirable results may occur.

If multiple Recordsets are retrieved, one table per Recordset, no data duplication is present. This may be performed efficiently by using a stored procedure to return multiple result sets in one round trip from a data source. However for multiple Recordsets, programming overheads are incurred in establishing relationships for navigational purposes.

Extra care is required when updating a Dataset to a database. The enforcement of constraints and cascading rules supports the integrity of a Dataset, however diligent management control is required in controlling and ensuring relationally connected Data

table changes are applied correctly to the data source. There is a possibility that the order in which changes are applied to a Dataset with multiple tables and relationships may not be applied correctly, i.e. in the same sequence when updating a database using a data adapter. That is, the update may produce incorrect results compromising the integrity of the database. ADO.NET enables Data table changes to be isolated by type of change. Consideration of change type and correct order of Data Table updates determined by parent/child relationships enables database updates to be performed in the correct order.

The ADO.NET Command Builder provides an alternative, but has limited control. There are no options to configure the update process. It is useful for single table access and updates.

To enable improved processing of Recordsets with relationships, Recordset Data Shaping (MS Data Shape provider) based on the Recordset may be used. It is related to the organisation, access, manipulation and display of Recordsets. It provides the ability to embed Recordsets in Recordsets, known as hierarchical Recordsets. It uses a special query syntax that uses the SHAPE clause with SELECT statements. It provides several alternative representations of hierarchies including Relation, Parameterised, Grouping and Hierarchy. These may be combined with each other in various ways to form a structure that facilitates drill-down search capability. A Recordset is updateable at any level in the hierarchy and is based on the use of optimistic or batch optimistic lock types. The shaped Recordset provides certain advantages such as, returning data that is less than from joined queries and applying single table updates, however submitting pending changes against multiple tables can be problematic. Data shaping has further limitations: the query syntax is not easy to follow, queries apply only to a single data source and filtering is difficult.

The Data Set may contain one or more Data Table objects that optionally have relationships between rows of the tables, may have primary or foreign key constraints, as well as cascading rules that operate against update and delete operations. The Dataset provides management support of its data through the use of collections e.g. Data Relation, Data Table Collection, Data Row Collection, Data Relations Collection and Data Column Collection. Data relationships may be created, deleted and modified dynamically. Data redundancy may be eliminated using individual tables instead of SELECT JOINS. Hierarchically related tables may be sourced from different data providers. ADO.NET does not require an additional data provider to support relationships between tables - the Data Relation object is implicitly provided.

The rows of data in a Data Table may be accessed, searched and modified and they may be the source of customized views using Data View objects. The Dataset may be filtered using the Select method to create an array of rows. Data View objects may be sorted, filtered, searched and bound to Web and Windows controls e.g. the Data View object may be bound to ASP.NET Web Server controls such as Data Grid, Repeater and Data List.

Compared to ADO, ADO.NET is more useful for representing hierarchical or network relationships between tables.

7. Dataset and Recordset Navigation:

A difference between ADO and ADO.NET is navigational access of Rows within a Recordset and a Data Set. Navigation within a Recordset is sequentially based - rows are accessed through the use of a cursor. The Recordset provides a client-side cursor or a server-side cursor. Using client-side cursors, updates are performed optimistically against a data source. Using server-side cursors, updates are performed either

pessimistically or optimistically. The use of a server-side pessimistic locking cursor provides greater control and visibility to data changes applied by other users at the data source.

There is no concept of an updateable server-side cursor in a Data Set. The Dataset is completely independent of any data source. The ADO.NET Data Reader however does maintain state on the server when reading the database data as a stream.

Navigation within a Data Table occurs through the use of indexing where Data Rows may be accessed randomly as well as sequentially. Navigation between Data Tables is supported with the Relations property of the Dataset, Child Relations and Parent Relations properties of the Data Table, and Get Child Rows and Get Parent Rows methods of the Data Row. Each of these relation properties contains a collection of Data Relations. Taken together, all these relational representations between tables of a Dataset provide for hierarchical and easy manipulation of parent/child relationships.

8. Conclusions and Recommendations:

This paper presented a comparative analysis of two Microsoft data access technology models: ADO and ADO.NET. The comparative analysis focused on relational aspects with particular attention to the ADO Recordset and the ADO.NET Dataset objects.

Providing a simplistic feature-versus-feature comparison is not totally realistic, however Table 1 (next page) attempts to reflect in a summarised form the differences by grouping similar features together.

Based on the table, it may be concluded that ADO is beneficial in instances where:

- ✓ scalability, that is demand on resources, is not an issue
- ✓ controlled access using persistent connections to the database is a priority
- ✓ controlled visibility to changes made to a database by other users is a priority

ADO would not be beneficial in instances where comprehensive support and ease of use is sought in processing hierarchically related Recordsets. ADO.NET provides benefits where

- ✓ there is a requirement for comprehensive management and control in batch updates
- ✓ ease of use is a priority in the update of disconnected data to a data source
- ✓ scalability is an issue
- ✓ the use of stored procedures in batch updating is important.

9. References:

1. Otey Michael, Conte Paul (2001): SQL Server 2000 Developers Guide, Osborne/McGraw-Hill.
2. Blehrud, C et al. (2000) Professional Windows DNA: Building Distributed Web Applications with VB, COM+, MSMQ, SOAP, and ASP, Wrox.
3. Bustos J. and Karli, W. (2002) Beginning .Net Web Services with VB.NET, Wrox.
4. Anderson Richard, Francis Brian, Homer Alex, Howard Rob, Sussman Dave, Watson Karli Professional ASP.NET 1.0, Wrox Press Ltd.
5. The Role of ADO in Universal Data Access, ADO Programmer's Guide
6. http://msdn.microsoft.com/library/en-us/ado270/htm/m_dconroleofadoinuda.asp,
7. MacDonald Rob: Serious ADO: Universal Data Access with Visual Basic, Apress.