



Data Providers for ADO.NET

Zen v15

Activate Your Data™

Copyright © 2023 Actian Corporation. All Rights Reserved.

This Documentation is for the end user's informational purposes only and may be subject to change or withdrawal by Actian Corporation ("Actian") at any time. This Documentation is the proprietary information of Actian and is protected by the copyright laws of the United States and international treaties. The software is furnished under a license agreement and may be used or copied only in accordance with the terms of that agreement. No part of this Documentation may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or for any purpose without the express written permission of Actian. To the extent permitted by applicable law, ACTIAN PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, AND ACTIAN DISCLAIMS ALL WARRANTIES AND CONDITIONS, WHETHER EXPRESS OR IMPLIED OR STATUTORY, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTY OF MERCHANTABILITY, OF FITNESS FOR A PARTICULAR PURPOSE, OR OF NON-INFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT WILL ACTIAN BE LIABLE TO THE END USER OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF ACTIAN IS EXPRESSLY ADVISED OF SUCH LOSS OR DAMAGE.

The manufacturer of this Documentation is Actian Corporation.

For government users, the Documentation is delivered with "Restricted Rights" as set forth in 48 C.F.R. Section 12.212, 48 C.F.R. Sections 52.227-19(c)(1) and (2) or DFARS Section 252.227-7013 or applicable successor provisions.

Actian, Actian DataCloud, Actian DataConnect, Actian X, Avalanche, Versant, PSQL, Actian Zen, Actian Director, Actian Vector, DataFlow, Ingres, OpenROAD, and Vectorwise are trademarks or registered trademarks of Actian Corporation and its subsidiaries. All other trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Contents

About This Document

ix

What Are the Zen Data Providers?	ix
Using This Guide	ix

Quick Start

1

ADO.NET Data Providers Installed with Zen	2
Supported .NET Framework Versions	2
Zen ADO.NET Data Providers Available with SDK Download	4
Defining Basic Connection Strings	6
Connecting to a Database	7
Example: Using the Provider-Specific Objects	7
Example: Using the Common Programming Model	9
Example: Using the Zen Common Assembly	9
Using the Zen ADO.NET Entity Framework Data Provider	11

Using the Data Providers

13

About the Data Providers	14
Using Connection Strings	15
Guidelines	15
Using the Zen Performance Tuning Wizard	16
Stored Procedures	17
Using IP Addresses	18
Transaction Support	19
Using Local Transactions	19
Thread Support	20
Unicode Support	21
Isolation Levels	22
SQL Escape Sequences	23
Event Handling	24
Error Handling	25
Using .NET Objects	26
Developing Applications for .NET	27

Advanced Features

29

Using Connection Pooling	30
Creating a Connection Pool	30

Adding Connections to a Pool	31
Removing Connections from a Pool	32
Handling Dead Connection in a Pool	32
Tracking Connection Pool Performance	33
Using Statement Caching	34
Enabling Statement Caching	34
Choosing a Statement Caching Strategy	35
Using Connection Failover	36
Using Client Load Balancing	38
Using Connection Retry	39
Configuring Connection Failover	40
Setting Security	42
Code Access Permissions	42
Security Attributes	42
Using Zen Bulk Load	43
Use Scenarios for Zen Bulk Load	43
Zen Common Assembly	44
Bulk Load Data File	45
Bulk Load Configuration File	45
Determining the Bulk Load Protocol	46
Character Set Conversions	46
External Overflow File	47
Bulk Copy Operations and Transactions	47
Using Diagnostic Features	49
Tracing Method Calls	49
PerfMon Support	50
Analyzing Performance With Connection Statistics	51
Enabling and Retrieving Statistical Items	52

The ADO.NET Data Providers 55

About Zen ADO.NET Data Providers	56
Namespace	56
Assembly Name	56
Using Connection Strings with the Zen ADO.NET Data Provider	57
Constructing a Connection String	57
Performance Considerations	58
Connection String Options that Affect Performance	58
Properties that Affect Performance	59
Data Types	60
Mapping Zen Data Types to .NET Framework Data Types	60

Mapping Parameter Data Types	62
Data Types Supported with Stream Objects	65
Using Streams as Input to Long Data Parameters	66
Parameter Markers	67
Parameter Arrays	68
Zen ADO.NET Core Data Providers	77
About Zen ADO.NET Core Data Providers	78
Creating an Application in Visual Studio Using Zen ADO.Net Core DLL	79
Creating a UWP Application in Visual Studio Using Zen ADO.Net Core Data Provider	81
ADO.NET Data Provider Features Missing in Zen ADO.NET Core Data Provider	84
Zen ADO.NET Entity Framework Data Providers	85
About Zen ADO.NET Entity Framework Data Providers	87
Namespace	87
Assembly Names	87
Configuring Entity Framework 6.1	88
Configuration File Registration	88
Code-Based Registration	89
Using Multiple Entity Framework Versions Against the Same Database	89
Using Connection Strings with the Zen ADO.NET Entity Framework Data Provider	90
Defining Connection String Values in Server Explorer	90
Changes in Default Values for Connection String Options	90
Code First and Model First Support	90
Handling Long Identifier Names	91
Using Code First Migrations with the ADO.NET Entity Framework	91
Using Enumerations with the ADO.NET Entity Framework	93
Mapping Data Types and Functions	94
Type Mapping for Database First	94
Type Mapping for Model First	96
Type Mapping for Code First	97
Mapping EDM Canonical Functions to Zen Functions	98
Extending Entity Framework Functionality	103
Enhancing Entity Framework Performance	104
Limiting the Size of XML Schema Files	104
Using Stored Procedures with the ADO.NET Entity Framework	105
Providing Functionality	105
Using Overloaded Stored Procedures	106
Using .NET Objects	107
Creating a Model	108

Upgrading Entity Framework 5 Applications to Entity Framework 6.1	114
For More Information	118
Zen ADO.NET Entity Framework Core Data Providers	119
About Zen ADO.NET Entity Framework Core Data Providers	120
Namespace	120
Assembly Names.	120
Configuring the Zen ADO.NET Entity Framework Core Data Provider	121
Using Connection Strings with the Zen ADO.NET Entity Framework Core Data Provider	122
Changes in Default Values for Connection String Options.	122
Code First Support	122
Handling Long Identifier Names.	122
Using Code First Migrations with the ADO.NET Entity Framework Core	123
Using Reverse Engineering (Scaffolding).	124
Type Mapping for Code First	125
Mapping EDM Canonical Functions to Zen Functions	127
Extending Entity Framework Functionality	132
Using Stored Procedures with the ADO.NET Entity Framework Core	133
Upgrading an Application from Entity Framework 6.x to Entity Framework Core	134
Limitations	135
For More Information	136
Using Zen Data Providers in Visual Studio	137
Adding Connections	138
Adding Connections in Server Explorer	138
Adding Connections with the Data Source Configuration Wizard	148
Using the Zen Performance Tuning Wizard	151
Using Provider-Specific Templates.	154
Creating a New Project	154
Adding a Template to an Existing Project.	155
Using the Zen Visual Studio Wizards	157
Creating Tables with the Add Table Wizard	157
Creating Views with the Add View Wizard.	161
Adding Components from the Toolbox.	164
Data Provider Integration Scenario	165
Using the Microsoft Enterprise Libraries	167
Data Access Application Blocks	168

When Should You Use the DAAB?	168
Should You Use Generic or Database-specific Classes?	168
Configuring the DAAB	169
Using the DAAB in Application Code	170
Logging Application Blocks	172
When Should You Use the LAB?	172
Configuring the LAB	172
Adding a New Logging Application Block Entry	173
Using the LAB in Application Code	173
Additional Resources	176

A. .NET Objects Supported 177

.NET Base Classes	178
Data Provider-Specific Classes	179
PsqlBulkCopy	180
PsqlBulkCopyColumnMapping	180
PsqlBulkCopyColumnMappingCollection	180
PsqlCommand Object	180
PsqlCommandBuilder Object	184
PsqlConnection Object	185
PsqlConnectionStringBuilder Object	187
PsqlCredential Object	198
PsqlDataAdapter Object	200
PsqlDataReader Object	200
PsqlError Object	201
PsqlErrorCollection Object	202
PsqlException Object	202
PsqlFactory Object	203
PsqlInfoMessageEventArgs Object	203
PsqlParameter Object	204
PsqlParameterCollection Object	205
PsqlTrace Object	206
PsqlTransaction Object	207
Zen Common Assembly Classes	208
CsvDataReader	208
CsvDataWriter	209
DbBulkCopy	210
DbBulkCopyColumnMapping	211
DbBulkCopyColumnMappingCollection	211

B. Getting Schema Information	213
Columns Returned by the GetSchemaTable Method	214
Retrieving Schema Metadata with the GetSchema Method	217
MetaDataCollections Schema Collections	217
DataSourceInformation Schema Collection	218
DataTypes Collection	219
ReservedWords Collection	221
Restrictions Collection	221
Additional Schema Collections	223
Columns Schema Collection	223
ForeignKeys Schema Collection	225
Indexes Schema Collection	227
PrimaryKeys Schema Collection	228
ProcedureParameters Schema Collection	229
Procedures Schema Collection	232
TablePrivileges Schema Collection	233
Tables Schema Collection	234
Views Schema Collection	234
C. SQL Escape Sequences for .NET	237
Date, Time, and Timestamp Escape Sequences	238
Scalar Functions	239
Outer Join Escape Sequences	240
D. Locking and Isolation Levels	241
Locking	242
Isolation Levels	243
Locking Modes and Levels	245
E. Designing .NET Applications for Performance Optimization	247
Retrieving Data	248
Retrieving Long Data	248
Reducing the Size of Data Retrieved	248
Using CommandBuilder Objects	249
Choosing the Right Data Type	249
Selecting .NET Objects and Methods	250
Using Parameter Markers as Arguments to Stored Procedures	250
Designing .NET Applications	251
Managing Connections	251

Opening and Closing Connections	251
Using Statement Caching	252
Using Commands Multiple Times	252
Using Native Managed Providers	253
Updating Data	254
Using the Disconnected DataSet	254
Synchronizing Changes Back to the Data Source	254
F. Using an .edmx File	255
Code Examples	256
G. Bulk Load Configuration Files	261
Sample Bulk Data Configuration File	262
XML Schema Definition for a Bulk Data Configuration File	263
H. IANA Code Page Mappings	265
I. Glossary	267
.NET Architecture	267
ADO.NET	267
ADO.NET Entity Framework	267
assembly	267
bulk load	268
client load balancing	268
code access security (CAS)	268
common language runtime (CLR)	268
connection failover	268
connection pooling	269
connection retry	269
Data Access Application Block (DAAB)	269
destination table	269
entity	269
global assembly cache (GAC)	269
isolation level	269
load balancing	270
locking level	270
Logging Application Block (LAB)	270
managed code	270
namespace	270
Performance Monitor	271

stream	271
schema collection	271
strong name	271
unmanaged code	271

About This Document

This documentation covers the Zen ADO.NET data provider, Zen ADO.NET Entity Framework data provider, and Zen ADO.NET Entity Framework Core data provider.

What Are the Zen Data Providers?

Zen ADO.NET data providers are managed data providers, built with 100% managed code. The data providers are native wire protocol providers, which means that the provider does not have to call out to unmanaged code, code outside of the .NET Framework, in the form of a database client unless your application enlists in Microsoft Distributed Transaction Coordinator (MS DTC) coordinated transactions.

The Zen ADO.NET data providers enable you to connect to Zen database engines. They work with both 32- and 64-bit .NET and are supported on all Zen supported Windows platforms.

See also [ADO.NET Data Providers Installed with Zen](#) and [Zen ADO.NET Data Providers Available with SDK Download](#).

Using This Guide

We assume you are familiar with your operating system and its commands, the definition of directories, and accessing a database through an end-user application.

This guide covers the following information:

- [Quick Start](#) provides information about connecting to a database with your .NET data provider.
- [Using the Data Providers](#) provides information about using .NET applications with the Zen data provider and provides information about developing .NET applications in the .NET environment.
- [Advanced Features](#) describes advanced features of the data providers, including connection pooling, statement caching, configuring security, and using Zen Bulk Load.
- [The ADO.NET Data Providers](#) describes connection string options, data types, and other information for the Zen Entity Framework data provider.
- [Zen ADO.NET Core Data Providers](#) describes how to create an application and a UWP application in Visual Studio using Zen ADO.NET Core DLL.

- [Zen ADO.NET Entity Framework Data Providers](#) describes features of the Zen ADO.NET Entity Framework data provider. It explains how to create a Entity Data Model for the Zen ADO.NET Entity Framework data provider.
- [Zen ADO.NET Entity Framework Core Data Providers](#) describes the Action Zen ADO.NET Entity Framework Core data provider and provides instructions on how to configure and use it.
- [Using Zen Data Providers in Visual Studio](#) describes how to use the Zen data provider and the Performance Wizard from within Visual Studio.
- [Using the Microsoft Enterprise Libraries](#) describes how to configure the Data Access Application Block and Logging Application Block, and use them in your application code.
- [.NET Objects Supported](#) provides the .NET public objects, properties, and methods supported by the Zen data provider.
- [Getting Schema Information](#) describes the schema collections supported by the Zen data provider.
- [SQL Escape Sequences for .NET](#) describes the scalar functions supported for the Zen data provider. Your data store may not support all of these functions.
- [Locking and Isolation Levels](#) discusses locking and isolation levels and how their settings can affect the data you retrieve.
- [Designing .NET Applications for Performance Optimization](#) provides recommendations for improving the performance of your applications by optimizing its code.
- [Using an .edmx File](#) explains the necessary changes to an .edmx file in order to provide Extended Entity Framework functionality to the EDM layer.
- [Bulk Load Configuration Files](#) provides samples of the files created during bulk load operations.
- [IANA Code Page Mappings](#) maps the most widely used IBM code pages to IANA code page names.
- [Glossary](#) defines terms used in this guide.

Note: This guide refers the reader to online links for more information about specific topics. Because it is the nature of web content to change frequently, we can guarantee only that the links referenced here were correct at the time of publishing.

Quick Start

The following basic information enables you to connect to a database using the Zen ADO.NET data providers immediately after their installation:

- [ADO.NET Data Providers Installed with Zen](#)
- [Zen ADO.NET Data Providers Available with SDK Download](#)
- [Defining Basic Connection Strings](#)
- [Connecting to a Database](#)
- [Using the Zen ADO.NET Entity Framework Data Provider](#)

To take full advantage of Zen ADO.NET data provider features, we recommend that you also read other Zen ADO.NET topics documented [here](#).

ADO.NET Data Providers Installed with Zen

This section describes the .NET Framework versions supported by the Zen ADO.NET data providers installed with Zen v15. For a description of Zen ADO.NET data providers for .NET Core and Entity Framework Core provided by downloading the SDK, see [Zen ADO.NET Data Providers Available with SDK Download](#).

Zen v15 provides two versions of the ADO.NET data providers: 4.4 and 4.5. All versions are installed by default with the database engine.

If you are using ADO.NET without customization, then code written for earlier versions of the .NET Framework and of the Zen data provider is compatible with Zen data provider 4.4 and 4.5.

Supported .NET Framework Versions

Zen ADO.NET Data Provider 4.4 and 4.5 add support for the combinations with Microsoft .NET Framework and Microsoft Entity Framework as shown in the following table. Each row of the table represents the compatible combinations of the supported versions of these three products.

Zen Data Provider for	Version	Namespace	Assembly File Name (Installed by Zen)	Microsoft .NET Framework	Microsoft Entity Framework
ADO.NET	4.4	Pervasive.Data.SqlClient	Pervasive.Data.SqlClient.dll	2.0, 3.0, 3.5, 3.5 SP1, 4.5, 4.5.1, 4.5.2, 4.6, 4.6.1, 4.6.2, 4.7, 4.7.1, 4.7.2, 4.8	—
ADO.NET	4.5	Pervasive.Data.SqlClient	Pervasive.Data.SqlClient.dll	2.0, 3.0, 3.5, 3.5 SP1, 4.5, 4.5.1, 4.5.2, 4.6, 4.6.1, 4.6.2, 4.7, 4.7.1, 4.7.2, 4.8	—
ADO.NET Entity Framework	4.4	Pervasive.Data.SqlClient.Entity	Pervasive.Data.SqlClient.Entity.dll	4.5, 4.5.1, 4.5.2, 4.6, 4.6.1, 4.6.2, 4.7, 4.7.1, 4.7.2, 4.8	6.1, 6.1.1, 6.1.2
ADO.NET Entity Framework	4.5	Pervasive.Data.SqlClient.Entity	Pervasive.Data.SqlClient.Entity.dll	4.5, 4.5.1, 4.5.2, 4.6, 4.6.1, 4.6.2, 4.7, 4.7.1, 4.7.2, 4.8	6.1, 6.1.1, 6.1.2

To use Zen ADO.NET Entity Framework Provider 4.4 or 4.5, your applications must target .NET Framework 4.5 or later.

All versions listed here apply to both 32- and 64-bit versions of .NET Framework.

Zen ADO.NET Entity Framework Provider 4.4 and 4.5 support Microsoft Entity Framework 6.1.

For more information on the Zen ADO.NET data providers, see the topic for each provider.

Zen ADO.NET Data Providers Available with SDK Download

Beyond the ADO.NET data providers installed with Zen v15, additional providers are available to support .NET Standard 2.0 compliant applications. These providers are listed in the following table and made available as NuGet packages included in a downloadable SDK at the [Actian website](#). As with the versions of providers installed with Zen, the following table lists two versions for compliant applications.

Zen Data Provider for	Version	Namespace	Assembly File Name	.NET	EF Core	NuGet Package (see note)	SDK Download
ADO.NET Core	4.4	Pervasive.Data.SqlClient	Pervasive.Data.SqlClientStd.dll	Core 2.0	—	Pervasive.Data.SqlClientStd.4.4.0.<build>.nupkg	Zen-SDK-AdoNetDataProvider4.4-NetStandard-Windows-noarch-<version>.zip
ADO.NET Core	4.5	Pervasive.Data.SqlClient	Pervasive.Data.SqlClientStd.dll	Core 2.1, 6.0, 7.0	—	Pervasive.Data.SqlClientStd.4.5.0.<build>.nupkg	Zen-SDK-AdoNetDataProvider4.5-NetStandard-Windows-noarch-<version>.zip
ADO.NET Entity Framework Core	4.4	Actian.EntityFrameworkCore.Zen	Actian.EntityFrameworkCore.Zen.dll	Core 2.0	3.1	Actian.EntityFrameworkCore.Zen.4.4.0.<build>.nupkg	Zen-SDK-AdoNetDataProvider4.4-NetStandard-Windows-noarch-<version>.zip
ADO.NET Entity Framework Core	4.5	Actian.EntityFrameworkCore.Zen	Actian.EntityFrameworkCore.Zen.dll	Core 2.1, 6.0	3.1, 6.0	Actian.EntityFrameworkCore.Zen.4.5.0.<build>.nupkg	Zen-SDK-AdoNetDataProvider4.5-NetStandard-Windows-noarch-<version>.zip

Note: The associated NuGet packages are available only in the downloadable SDK listed in the table.

Zen ADO.NET Entity Framework Core data providers require the Zen ADO.NET Core data provider of the same version to be added to the application project.

For more information on the Zen ADO.NET Core data providers, see the topic for each provider.

Defining Basic Connection Strings

The data provider uses a connection string to provide information needed to connect to a specific database server. The connection information is defined by connection string options.

The Zen ADO.NET Entity Framework data provider can specify an existing connection in the Entity Framework Wizard, or can define a new connection. The Zen ADO.NET Entity Framework uses information contained in connection strings to connect to the underlying Zen ADO.NET data provider that supports the Entity Framework. The connection strings also contain information about the required model and mapping files. The data provider uses the connection string when accessing model and mapping metadata and connecting to the data source.

The connection string options have the form:

```
"option name=value"
```

Each connection string option value pair is separated by a semicolon. For example,

```
"Server DSN=DEMADATA;UID=test;PWD=test;Host=localhost"
```

See [Connection String Properties](#) for details about the supported connection string options.

Notes

- The spaces in the option names are optional.
- All connection string option names are case-insensitive. For example, User ID is the same as user id. However, the values of some options, such as User ID and Password, might be case-sensitive.
- If the connection string does not specify a port number, the data provider uses 1583, the default port number.

Minimum Required Connection String Options

The following table gives the name and description for each option required for a minimum connection to a Zen server.

Option	Description
Server DSN	Specifies the name of the data source on the server to which you want to connect, for example, DEMADATA.
Host	Specifies the name or the IP address of the Zen server to which you want to connect. For example, you can specify a server name such as Accountingserver or an IP address such as 199.226.22.34 (IPv4) or 1234:5678:0000:0000:0000:0000:9abc:def0 (IPv6). The initial default value is localhost.

Connecting to a Database

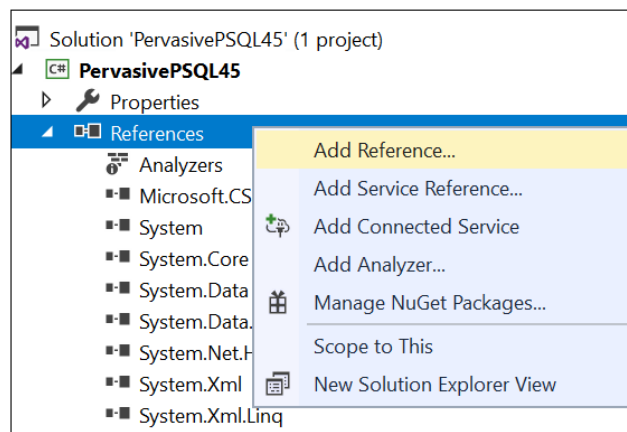
Once your data provider is installed, you can connect from your application to your database with a connection string. See [Connection String Properties](#) for a list of the connection string options.

Note: If your application uses the Zen ADO.NET Entity Framework, you can use the Entity Data Model Wizard to create a new connection or use an existing connection. See [Creating a Model](#) for more information.

Example: Using the Provider-Specific Objects


The following example uses the provider-specific objects to connect to a database using the Zen ADO.NET data provider from an application developed in Visual Studio using C#.

1. In the Solution Explorer, right-click **References** and then select **Add Reference**.

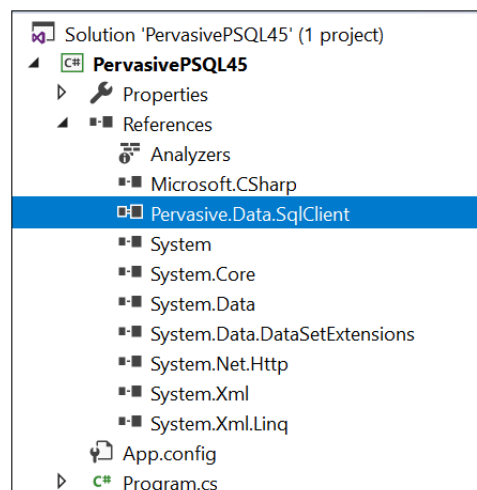


2. In the **Reference Manager** wizard, click the **Browse** button and navigate to the folder that contains the Zen data provider assembly.

3. Select **Pervasive.Data.SqlClient.dll** and click **Add**. The **Browse** tab of the Reference Manager wizard lists the Zen data provider assembly in the **Recent** items.



4. Select it and click **OK**. The Solution Explorer now includes the Zen data provider.



5. Add the Zen data provider's namespace to the beginning of your application, as shown in the following C# code fragment:

```
// Access Zen
using System.Data;
using System.Data.Common;
using Pervasive.Data.SqlClient;
```

6. Add exception handling code and the connection information for your server:

```
PsqlConnection DBConn = new PsqlConnection("Server DSN=DEMODATA;Host=localhost");
try
{
    // Open the connection
    DBConn.Open();
    Console.WriteLine("Connection Successful!")
}
```

```
}  
catch (PsqlException ex)  
{  
    // Connection failed  
    writer.WriteLine(ex.Message);  
}
```

7. Close the connection.

```
// Close the connection  
DBConn.Close();
```

Example: Using the Common Programming Model

The following example illustrates connecting to a Zen database from an application developed in Visual Studio using C# and the Common Programming Model.

1. Check the beginning of your application. Ensure that the ADO.NET namespaces are present.

```
// Access Zen using factory  
using System.Data;  
using System.Data.Common;
```

2. Add the connection information of your server and exception handling code and close the connection.

```
DbProviderFactory  
factory=DbProviderFactories("Pervasive.Data.SqlClient");  
DbConnection Conn = factory.createConnection();  
Conn.CommandText = "Server DSN=DEMODATA;Host=localhost;";  
try  
{  
    Conn.Open();  
    Console.WriteLine("Connection successful!");  
}  
catch (Exception ex)  
{  
    // Connection failed  
    Console.WriteLine(ex.Message);  
}  
// Close the connection  
Conn.Close();
```

Example: Using the Zen Common Assembly

You can optionally include the Zen Common Assembly if you want to use features such as Zen Bulk Load in an application that conforms to the Common Programming Model. See [Using Zen Bulk Load](#) for information about how to use Zen Bulk Load with your application.

The following example illustrates how to use the Zen Common Assembly in an application developed in Visual Studio using C# and the Common Programming Model.

-
1. Check the beginning of your application. Ensure the .NET Framework and Zen data provider namespaces are present.

```
// Access Zen using factory
using System.Data;
using System.Data.Common;
using Pervasive.Data.Common;
```

2. Add the connection information of your server and exception handling code and close the connection.

```
// This code does a bulk copy operation from
// one database to another
DbProviderFactory Factory = DbProviderFactories.GetFactory("Pervasive.Data.SqlClient");
DbConnection sourceConnection = Factory.CreateConnection();
sourceConnection.ConnectionString = "Host=localhost;Server DSN=DEMODATA;";

sourceConnection.Open();

DbCommand command = sourceConnection.CreateCommand();
command.CommandText = "SELECT * FROM test";
DbDataReader reader = command.ExecuteReader();

DbConnection destinationConnection = Factory.CreateConnection();
destinationConnection.ConnectionString =
    "Host= ntssl2003b;Server DSN=DEMODATA";
destinationConnection.Open();

DbBulkCopy bulkCopy = new DbBulkCopy(destinationConnection);
bulkCopy.DestinationTableName = "test";
try
{
    bulkCopy.WriteToServer(reader);
} //end try
catch (DbException ex)
{
    Console.WriteLine( ex.Message );
} //end catch
finally
{
    reader.Close();
    MessageBox.Show("done");
} //end finally
```

Using the Zen ADO.NET Entity Framework Data Provider

The Entity Data Model wizard asks questions that help you to define the components in your Entity Data Model (EDM). The wizard then creates a model of your data in Visual Studio, and automatically sets values for the components in the model. See [Using an .edmx File](#) for information about using the wizard to create an EDM.

Alternatively, you can use other tools in Visual Studio to define values and connection strings manually.

Provider is an attribute of the Schema element in the storage model file of an EDM. The storage model file is written in the store schema definition language (SSDL).

The Entity Data Model wizard assigns the value when you select the Zen ADO.NET Entity Framework data provider. If you choose to manually define an Entity Data Model, assign the string `Pervasive.Data.SqlClient` to the Provider attribute of the Schema element, as shown in the following example:

```
<Schema Namespace="AdventureWorksModel.Store" Alias="Self" Provider="Pervasive.Data.SqlClient"
ProviderManifestToken="Zen" xmlns:store="http://schemas.microsoft.com/ado/2007/12/edm/
EntityStoreSchemaGenerator" xmlns="http://schemas.microsoft.com/ado/2006/04/edm/ssdl">
```


Using the Data Providers

The Zen data providers provide data access to any .NET-enabled application or application server. The data providers delivers high-performance point-to-point and *n*-tier access to industry-leading data stores across the Internet and intranets. Because they are optimized for the .NET environment, the data providers allow you to incorporate .NET technology and extend the functionality and performance of your existing system.

See [Advanced Features](#) for information on advanced features such as connection pooling, statement caching, configuring security, Zen Bulk Load, and diagnostic support.

See [The ADO.NET Data Providers](#) for information about using the Zen ADO.NET data provider in the standard Zen ADO.NET environment.

See [Zen ADO.NET Entity Framework Data Providers](#) for information about using the data provider with the Zen ADO.NET Entity Framework.

About the Data Providers

Zen data providers are built with 100% managed code, so they can run and connect to the database entirely within the common language runtime (CLR).

Code that runs in the native operating system, such as client libraries and COM components, is called unmanaged code. You can mix managed and unmanaged code within a single application. However, unmanaged code reaches outside the CLR, which means that it effectively increases complexity, reduces performance, and opens possible security risks.

Using Connection Strings

You can define the behavior of a connection using a connection string or the properties of the `PsqlConnection` object.

However, values set in the connection string cannot be changed by the connection properties.

The basic format of a connection string includes a series of keyword/value pairs separated by semicolons. The following example shows the keywords and values for a simple connection string for the Zen data provider:

```
"Server DSN=SERVERDEMO;Host=localhost"
```

Guidelines

Use the following guidelines when specifying a connection string:

- The spaces in the connection string option names are required.
- All connection string option names are case-insensitive. For example, `Password` is the same as `password`. However, the values of options such as `User ID` and `Password` may be case-sensitive.
- To include values that contain a semicolon, single quote, or double quotes, enclose the value in double quotes. If the value contains both a semicolon and double quotes, use single quotes to enclose the value.
- You can also use single quotes when the value starts with a double quote. Conversely, double quotes can be used if the value starts with a single quote. If the value contains both single quotes and double quotes, the character used to enclose the value must be doubled every time it occurs within the value.
- To include leading or trailing spaces in the string value, the value must be enclosed in either single quotes or double quotes. Any leading or trailing spaces around integer, Boolean, or enumerated values are ignored, even if enclosed in single or double quotes. However, spaces within a string literal keyword or value are preserved. Single or double quotes can be used within a connection string without using delimiters (for example, `Data Source= my'Server` or `Data Source= my"Server`), unless it is the first or last character in the value.
- Special characters can be used in the value of the connection string option. To escape special characters, surround the value in single or double quotes.
- The Equals character (`=`) can also be repeated within the connection string. For example:

```
Initialization String=update mytable set coll == 'foo'"
```

-
- If the connection string contains invalid connection string options, the connection attempt returns an error. For example, an error is returned if you specify a value for Load Balancing when Alternate Servers has not been defined.
 - If the connection string contains duplicated connection string options, the data provider uses the connection string option that appears last in the connection string. For example, Connection Timeout appears twice in the following connection string, with different values. The data provider uses the second value and waits 35 seconds before terminating an attempted connection:

```
"Server DSN=SERVERDEMO;Host=localhost;Connection Timeout=15;Min Pool Size=50;Connection Timeout=35"
```

See [Connection String Properties](#) for a list of the supported connection string options.

Using the Zen Performance Tuning Wizard

You can use the Performance Wizard to select the optimal connection string options for both the Zen ADO.NET data provider or the Zen ADO.NET Entity Framework data provider.

See [Using the Zen Performance Tuning Wizard](#) for more information.

Stored Procedures

To enable stored procedures in the application, do the following:

- Set the CommandText property in the PsqlCommand object to the stored procedure name.

```
MyCommand.CommandText = "GetEmpSalary";
```

- Set the CommandType property in the PsqlCommand object to StoredProcedure.

```
MyCommand.CommandType = CommandType.StoredProcedure;
```

- Specify parameter arguments, if needed. The application should add the parameters to the parameter collection of the PsqlCommand object in the order of the arguments to the stored procedure. The application does not need to specify the parameter markers in the CommandText property of the PsqlCommand object.

To retrieve the return value from a stored procedure, the application should add an extra parameter to the parameter collection for the PsqlCommand object. This parameter's ParameterDirection property should be set to ParameterDirection.ReturnValue. The return value parameter can be anywhere in the parameter collection because it does not correspond to a specific parameter marker in the Text property of the PsqlCommand object.

If the stored procedure does not produce a return value, parameters bound with the ParameterDirection property as ReturnValue are ignored.

If the stored procedure returns a ReturnValue from the database and the application has not bound a parameter for it, the data provider discards the value.

Note for Zen ADO.NET Entity Framework Users: The PsqlConnection object includes properties and methods that provide enhanced statistics functionality. The methods and properties are standard in the Zen ADO.NET data provider but are not available at the Zen ADO.NET Entity Framework layer. Instead, the Zen ADO.NET Entity Framework data provider exposes the same functionality through "pseudo" stored procedures. See [Using Stored Procedures with the ADO.NET Entity Framework](#) for more information.

Using IP Addresses

The data providers support Internet Protocol (IP) addresses in IPv4 and IPv6 formats. If your network supports named servers, the server name specified in the data source can resolve to an IPv4 or an IPv6 address.

The EnableIPv6 connection string option, when set to True, allows a client with IPv6 protocol installed to connect to the server using either an IPv4 address or an IPv6 address. For more information about IPv6 formats, see [IPv6](#) in *Getting Started with Zen*.

Transaction Support

Zen data providers use only 100% managed code to support the transactions, which are implemented entirely within the .NET Framework.

Using Local Transactions

Local transactions use the internal transaction manager of the underlying database.

The application creates a `PsqlTransaction` object by calling `BeginTransaction` on the `PsqlConnection` object. Subsequent operations, such as committing or aborting the transaction, are performed on the `PsqlTransaction` object.

Thread Support

The `PsqlConnection` object is thread-safe. Multiple `PsqlCommand` objects, each accessed on a separate thread, can simultaneously use a single connection.

Accessing other public and data provider-specific objects simultaneously on separate threads is not thread-safe.

Unicode Support

The data provider supports Unicode as specified in the .NET Framework SDK. Effectively, this means that the data provider uses Unicode UTF-16 encoding to represent characters.

The data provider converts UTF-16 characters to the format used by the database, and returns .NET Framework strings to the application. For example, if a Zen database code page is in extended ASCII format, the data provider uses extended ASCII to represent characters sent to the database. The data provider then converts the extended ASCII characters returned before sending them back to the application.

For more information about the .NET Framework implementation of Unicode and international characters, refer to the Microsoft .NET Framework SDK documentation.

Isolation Levels

Zen supports the ReadCommitted and Serializable isolation levels. It supports record-level locking. See [Locking and Isolation Levels](#) for details.

SQL Escape Sequences

See [SQL Escape Sequences for .NET](#) for information about the SQL escape sequences supported by the Zen data provider.

Event Handling

The event handler receives an argument of type `PsqlInfoMessageEventArgs`, which contains data relevant to an event. See [PsqlInfoMessageEventArgs Object](#) for more information.

This event is defined as:

```
public event PsqlInfoMessageEventHandler InfoMessage;
```

Clients that want to process warnings and informational messages sent by the database server should create an `PsqlInfoMessageEventHandler` delegate to listen to this event.

You can use these events to capture failures that can occur when creating packages, stored procedures, or stored functions, which all create commands. If Zen encounters errors when compiling a command created by a package, stored procedure, or stored function, the object is created, even though it is not valid. An event will be sent, indicating the failure.

The following code fragment defines a delegate that represents the method that handles the `InfoMessage` event of a `PsqlConnection` object:

```
[Serializable]
public delegate void PsqlInfoMessageEventHandler(
    object sender
    PsqlInfoMessageEventArgs e
);
```

where *sender* is the object that generated the event and *e* is an `PsqlInfoMessageEventArgs` object that describes the warning. For more information on Events and Delegates, refer to the .NET Framework SDK documentation.

Error Handling

The `PsqlError` object collects information relevant to errors and warnings generated by the Zen server. See [PsqlError Object](#) for more information.

The `PsqlException` object is created and thrown when the Zen server returns an error. Exceptions generated by the data provider are returned as standard run time exceptions. See [PsqlException Object](#) for more information.

Using .NET Objects

The data provider supports the .NET public objects, exposing them as sealed objects.

See [.NET Objects Supported](#) for more information.

Developing Applications for .NET

Developers of data consumer applications must be familiar with the Microsoft .NET specification and object-oriented programming techniques.

Microsoft also provides extensive information about ADO.NET online, including the following articles:

- *Upgrading to Microsoft .NET: ADO.NET for the ADO Programmer*
<http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dndotnet/html/upgradingtodotnet.asp>
- *Using .NET Framework Data Providers to Access Data*
[http://msdn2.microsoft.com/en-us/library/s7ee2dwt\(vs.71\).aspx](http://msdn2.microsoft.com/en-us/library/s7ee2dwt(vs.71).aspx)
- *Generic Coding with the ADO.NET 2.0 Base Classes and Factories*
<http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnvs05/html/vsgenerics.asp>
- *Security Policy Best Practices*
[http://msdn.microsoft.com/en-us/library/sa4se9bc\(v=vs.100\).aspx](http://msdn.microsoft.com/en-us/library/sa4se9bc(v=vs.100).aspx)
- *Writing Serviced Components*
[http://msdn2.microsoft.com/en-us/library/3x7357ez\(vs.71\).aspx](http://msdn2.microsoft.com/en-us/library/3x7357ez(vs.71).aspx)
- *DataSets, DataTables, and DataViews*
[http://msdn.microsoft.com/en-us/library/ss7fbaez\(vs.110\).aspx](http://msdn.microsoft.com/en-us/library/ss7fbaez(vs.110).aspx)
- *Using XML in a DataSet*
[http://msdn.microsoft.com/en-us/library/84sxtbxh\(v=vs.110\).aspx](http://msdn.microsoft.com/en-us/library/84sxtbxh(v=vs.110).aspx)

Advanced Features

The following topics describe advanced features of the data provider:

- [Using Connection Pooling](#)
- [Using Statement Caching](#)
- [Using Connection Failover](#)
- [Using Client Load Balancing](#)
- [Using Connection Retry](#)
- [Setting Security](#)
- [Using Zen Bulk Load](#)
- [Using Diagnostic Features](#)

Using Connection Pooling

Connection pooling allows you to reuse connections rather than creating a new one every time the data provider needs to establish a connection to the underlying database. The data provider automatically enables connection pooling for your .NET client application.

You can control connection pooling behavior by using connection string options. For example, you can define the number of connection pools, the number of connections in a pool, and the number of seconds before a connection is discarded.

Connection pooling in ADO.NET is not provided by the .NET Framework. It must be implemented in the ADO.NET data provider itself.

Creating a Connection Pool

Each connection pool is associated with a specific connection string. By default, when the first connection with a unique connection string connects to the database, the connection pool is created. The pool is populated with connections up to the minimum pool size. Additional connections can be added until the pool reaches the maximum pool size.

The pool remains active as long as any connections remain open, either in the pool or used by an application with a reference to a Connection object that has an open connection.

If a new connection is opened and the connection string does not match an existing pool, a new pool must be created. By using the same connection string, you can enhance the performance and scalability of your application.

In the following C# code fragments, three new PsqlConnection objects are created, but only two connection pools are required to manage them. Note that the first and second connection strings differ only by the values assigned for User ID and Password, and by the value of the Min Pool Size option.

```
DbProviderFactory Factory = DbProviderFactories.GetFactory("Pervasive.Data.SqlClient");
DbConnection conn1 = Factory.CreateConnection();
conn1.ConnectionString = "Server DSN=DEMOMDATA;User ID=test;
Password = test; Host = localhost;MinPoolSize=5 ";
conn1.Open();
// Pool A is created.
DbConnection conn2 = Factory.CreateConnection();
conn2.ConnectionString = "Server DSN=DEMOMDATA2;User ID=lucy;
Password = quake; Host = localhost;MinPoolSize=10 ";
conn2.Open();
// Pool B is created because the connection strings differ.
DbConnection conn3 = Factory.CreateConnection();
conn3.ConnectionString = "Server DSN=DEMOMDATA;User ID=test;
Password = test; Host = localhost;MinPoolSize=5 ";
conn3.Open();
// conn3 goes into Pool A with conn1.
```

Adding Connections to a Pool

A connection pool is created in the process of creating each unique connection string that an application uses. When a pool is created, it is populated with enough connections to satisfy the minimum pool size requirement, set by the Min Pool Size connection string option. If an application is using more connections than Min Pool Size, the data provider allocates additional connections to the pool up to the value of the Max Pool Size connection string option, which sets the maximum number of connections in the pool.

When a Connection object is requested by the application calling the Connection.Open(...) method, the connection is obtained from the pool, if a usable connection is available. A usable connection is defined as a connection that is not currently in use by another valid Connection object, has a matching distributed transaction context (if applicable), and has a valid link to the server.

If the maximum pool size has been reached and no usable connection is available, the request is queued in the data provider. The data provider waits for the value of the Connection Timeout connection string option for a usable connection to return to the application. If this time period expires and no connection has become available, then the data provider returns an error to the application.

You can allow the data provider to create more connections than the specified maximum pool size without affecting the number of connections pooled. This may be useful, for example, to handle occasional spikes in connection requests. By setting the Max Pool Size Behavior connection string option to SoftCap, the number of connections created can exceed the value set for Max Pool Size, but the number of connections pooled does not. When the maximum connections for the pool are in use, the data provider creates a new connection. If a connection is returned and the pool contains idle connections, the pooling mechanism selects a connection to be discarded so that the connection pool never exceeds the Max Pool Size. If Max Pool Size Behavior is set to HardCap, the number of connections created does not exceed the value set for Max Pool Size.

Important: Closing the connection using the Close() or Dispose() method of the PsqlConnection object adds or returns the connection to the pool. When the application uses the Close() method, the connection string settings remain as they were before the Open() method was called. If you use the Dispose method to close the connection, the connection string settings are cleared, and the default settings are restored.

Removing Connections from a Pool

A connection is removed from a connection pool when it either exceeds its lifetime as determined by the Load Balance Timeout connection string option, or when a new connection that has a matching connection string is initiated by the application (`PsqlConnection.Open()` is called).

Before returning a connection from the connection pool to an application, the Pool Manager checks to see if the connection has been closed at the server. If the connection is no longer valid, the Pool Manager discards it, and returns another connection from the pool, if one is available and valid.

You can control the order in which a connection is removed from the connection pool for reuse, based on how frequently or how recently the connection has been used, with the Connection Pool Behavior connection string option. For a balanced use of connections, use the `LeastFrequentlyUsed` or `LeastRecentlyUsed` values. Alternatively, for applications that perform better when they use the same connection every time, you can use the `MostFrequentlyUsed` or `MostRecentlyUsed` values.

The `ClearPool` and `ClearAllPools` methods of the Connection object remove all connections from connection pools. *ClearPool* clears the connection pool associated with a specific connection. In contrast, *ClearAllPools* clears all of the connection pools used by the data provider. Connections that are in use when the method is called are discarded when they are closed.

Note: By default, if discarding an invalid connection causes the number of connections to drop below the number specified in the Min Pool Size attribute, a new connection will not be created until an application needs one.

Handling Dead Connection in a Pool

What happens when an idle connection loses its physical connection to the database? For example, suppose the database server is rebooted or the network experiences a temporary interruption. An application that attempts to connect using an existing Connection object from a pool could receive errors because the physical connection to the database has been lost.

The Control Center handles this situation transparently to the user. The application does not receive any errors on the `Connection.Open()` attempt because the data provider simply returns a connection from a connection pool. The first time the Connection object is used to execute a SQL statement (for example, through the `Execute` method on the Command object), the data provider detects that the physical connection to the server has been lost and attempts to reconnect to the server *before* executing the SQL statement. If the data provider can reconnect to the server, the result of the SQL execution is returned to the application; no errors are returned to the application. The data provider uses the connection failover options, if enabled, when attempting this seamless

reconnection. See [Using Connection Failover](#) for information about configuring the data provider to connect to a backup server when the primary server is not available.

Note: Because the data provider can attempt to reconnect to the database server when executing SQL statements, connection errors can be returned to the application when a statement is executed. If the data provider cannot reconnect to the server (for example, because the server is still down), the execution method throws an error indicating that the reconnect attempt failed, along with specifics about the reason the connection failed.

This technique for handling dead connections in connection pools allows for the maximum performance out of the connection pooling mechanism. Some data providers periodically ping the server with a dummy SQL statement while the connections sit idle. Other data providers ping the server when the application requests the use of the connection from the connection pool. Both of these approaches add round trips to the database server and ultimately slow down the application during normal operation of the application is occurring.

Tracking Connection Pool Performance

The data providers install a set of PerfMon counters that let you tune and debug applications that use the data provider. See [PerfMon Support](#) for information about using the PerfMon counters.

Using Statement Caching

A statement cache is a group of prepared statements or instances of Command objects that can be reused by an application. Using statement caching can improve application performance because the actions on the prepared statement are performed once even though the statement is reused multiple times over an application's lifetime. You can analyze the effectiveness of the statements in the cache (see [Analyzing Performance With Connection Statistics](#)).

A statement cache is owned by a physical connection. After being executed, a prepared statement is placed in the statement cache and remains there until the connection is closed.

Statement caching can be used across multiple data sources and can be used beneath abstraction technologies such as the Microsoft Enterprise Libraries with the Data Access Application Blocks.

Enabling Statement Caching

By default, statement caching is not enabled. To enable statement caching for existing applications, set the Statement Cache Mode connection string option to Auto. In this case, all statements are eligible to be placed in the statement cache.

You can also configure statement caching so that only statements that you explicitly mark to be cached are placed in the statement cache. To do this, set the StatementCacheBehavior property of the statement's Command object to Cache and set the Statement Cache Mode connection string option to ExplicitOnly.

The following table summarizes the statement caching settings and their effects.

Behavior	StatementCacheBehavior	Statement Cache Mode
Explicitly add the statement to the statement cache.	Cache	ExplicitOnly (the default)
Add the statement to the statement cache. If necessary, the statement is removed to make room for a statement marked Cache.	Implicit (the default)	Auto
Specifically exclude the statement from the statement cache.	DoNotCache	Auto or ExplicitOnly

Choosing a Statement Caching Strategy

Statement caching provides performance gains for applications that reuse prepared statements multiple times over the lifetime of an application. You set the size of the statement cache with the Max Statement Cache Size connection string option. If space in the statement cache is limited, do not cache prepared statements that are used only once.


Caching all of the prepared statements that an application uses might appear to offer the best performance. However, this approach may come at a cost of database memory if you implement statement caching with connection pooling. In this case, each pooled connection has its own statement cache that may contain all of the prepared statements used by the application. All of these pooled prepared statements are also maintained in the database's memory.

Using Connection Failover

Connection failover allows an application to connect to an alternate, or backup, database server if the primary database is unavailable, for example, because of a hardware failure or traffic overload. Connection failover ensures that the data on which your critical .NET applications depend is always available.

You can customize the data provider for connection failover by configuring a list of alternate databases that are tried if the primary server is not accepting connections. Connection attempts continue until a connection is successfully established or until all of the alternate databases have been tried the specified number of times.

For example, the following figure shows an environment with multiple database servers. Database Server A is designated as the primary database server, Database Server B is the first alternate server, and Database Server C is the second alternate server.



First, the application attempts to connect to the primary database, Database Server A (1). If connection failover is enabled and Database Server A fails to accept the connection, the application attempts to connect to Database Server B (2). If that connection attempt also fails, the application attempts to connect to Database Server C (3).

In this scenario, it is probable that at least one connection attempt would succeed, but if no connection attempt succeeds, the data provider can retry the primary server and each alternate database for a specified number of attempts. You can specify the number of attempts that are made through the *connection retry* feature. You can also specify the number of seconds of delay, if any, between attempts through the *connection delay* feature. For more information about connection retry, see [Using Connection Retry](#).

The data provider fails over to the next alternate server only if it cannot establish communication with the current alternate server. If the data provider successfully establishes communication with a database and the database rejects the connection request because, for example, the login information is invalid, then the data provider generates an exception and does not try to connect to


the next database in the list. It is assumed that each alternate server is a mirror of the primary and that all authentication parameters and other related information are the same.

Connection failover provides protection for new connections only and does not preserve states for transactions or queries. For details on configuring connection failover for your data provider, see [Configuring Connection Failover](#).

Using Client Load Balancing

Client load balancing works with connection failover to distribute new connections in your environment so that no one server is overwhelmed with connection requests. When both connection failover and client load balancing are enabled, the order in which primary and alternate databases are tried is random.

For example, suppose that client load balancing is enabled:



First, Database Server B is tried (1). Then, Database Server C may be tried (2), followed by a connection attempt to Database Server A (3); subsequent connection attempts use this same sequence. In contrast, if client load balancing was not enabled in this scenario, each database would be tried in sequential order, primary server first, then alternate servers based on their entry order in the alternate servers list.

For details on configuring client and load balancing for your data provider, see [Configuring Connection Failover](#).

Using Connection Retry

Connection retry defines the number of times that the data provider attempts to connect to the primary, and, if configured, alternate database servers after the first unsuccessful connection attempt. Connection retry can be an important strategy for system recovery. For example, suppose you have a power failure scenario in which both the client and the server fail. When the power is restored and all computers are restarted, the client may be ready to attempt a connection before the server has completed its startup routines. If connection retry is enabled, the client application can continue to retry the connection until a connection is successfully accepted by the server.

Connection retry can be used in environments that only have one server or can be used as a complementary feature in connection failover scenarios with multiple servers.

Using connection string options, you can specify the number of times the data provider attempts to connect and the time in seconds between connection attempts. For details on configuring connection retry, see [Configuring Connection Failover](#).

Configuring Connection Failover

Connection failover allows an application to connect to an alternate, or backup, database server if the primary database server is unavailable, for example, because of a hardware failure or traffic overload.

See [Using Connection Failover](#) for more information about connection failover.

To configure connection failover to another server, you must specify a list of alternate database servers that are tried at connection time if the primary server is not accepting connections. To do this, use the Alternate Servers connection string option. Connection attempts continue until a connection is successfully established or until all the databases in the list have been tried once (the default).

Optionally, you can specify the following additional connection failover features:

- The number of times the data provider attempts to connect to the primary and alternate servers after the initial connection attempt. By default, the data provider does not retry. To set this feature, use the Connection Retry Count connection string option.
- The wait interval, in seconds, used between attempts to connect to the primary and alternate servers. The default interval is 3 seconds. To set this feature, use the Connection Retry Delay connection option.
- Whether the data provider will use load balancing in its attempts to connect to primary and alternate servers. If load balancing is enabled, the data provider uses a random pattern instead of a sequential pattern in its attempts to connect. The default value is not to use load balancing. To set this feature, use the Load Balancing connection string option.

You use a connection string to direct the data provider to use connection failover. See [Using Connection Strings](#).

The following C# code fragment includes a connection string that configures the data provider to use connection failover in conjunction with all of its optional features – load balancing, connection retry, and connection retry delay:

```
Conn = new PsqlConnection Conn = new PsqlConnection();
Conn = new PsqlConnection("Host=myServer;User ID=test;Password=secret;
Server DSN=SERVERDEMO;Alternate Servers="Host=AcctServer, Host=AcctServer2";
Connection Retry Count=4;Connection Retry Delay=5;Load Balancing=true;
Connection Timeout=60")
```

Specifically, this connection string configures the data provider to use two alternate servers as connection failover servers, to attempt to connect four additional times if the initial attempt fails, to wait five seconds between attempts, and to try the primary and alternate servers in a random

order. Each connection attempt lasts for 60 seconds, and uses the same random order that was established on the first retry.

Setting Security

The data provider supports Encrypted Network Communications, also known as wire encryption, on connections. By default, the data provider reflects the server setting.

The level of encryption allowed by the data provider depends on the encryption module used. With the default encryption module, the data provider supports 40-, 56-, and 128-bit encryption.

Data encryption may adversely affect performance because of the additional overhead, mainly CPU usage, required to encrypt and decrypt data.

In addition to encryption, the Control Center implements security through the security permissions defined by the .NET Framework.

Code Access Permissions

The data provider requires the FullTrust permission to be set in order to load and run. This requirement is due to underlying classes in System.Data that demand FullTrust for inheritance. All ADO.NET data providers require these classes to implement a DataAdapter.

Security Attributes

The data provider is marked with the AllowPartiallyTrustedCallers attribute.

Using Zen Bulk Load

Zen Bulk Load offers a one-stop approach for all of your bulk load needs, with a simple, consistent way to do bulk load operations for Zen and for all of the DataDirect Connect products that support this bulk load feature. This means that you can write your bulk load applications using the standards-based API bulk interfaces, and then, just plug in the database data providers or drivers to do the work for you.


Suppose you need to load data into Zen, Oracle, DB2, and Sybase. In the past, you probably had to use a proprietary tool from each database vendor for bulk load operations, or write your own tool. Now, because of the interoperability built into Zen Bulk Load, your task is much easier. Another advantage is that Zen Bulk Load uses 100% managed code, and requires no underlying utilities or libraries from other vendors.

Bulk load operations between dissimilar data stores are accomplished by persisting the results of the query in a comma-separated value (CSV) format file, a bulk load data file. The file can be used between the Control Center and any DataDirect Connect for Zen ADO.NET data providers that support bulk load. In addition, the bulk load data file can be used with any DataDirect Connect product driver or data provider that supports the Bulk load functionality. For example, the CSV file generated by the Zen data provider can be used by a DataDirect Connect for ODBC driver that supports bulk load.

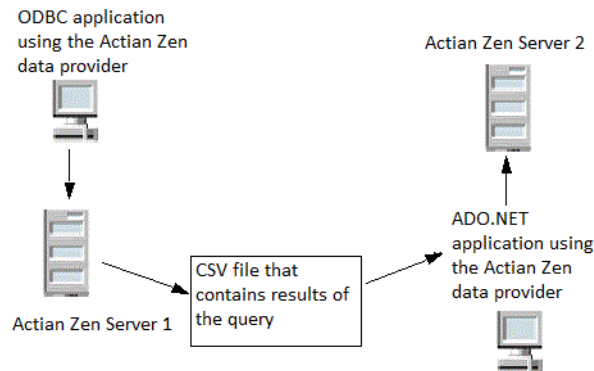
Use Scenarios for Zen Bulk Load

You can use Zen Bulk Load with the Control Center in two ways:

- Upgrade to a new Zen version and bulk-copy data from the old Zen data source to the new one, as shown in the following figure:



- Export data from a database and migrate the results to a Zen database. The following figure shows an ODBC environment copying data to an Zen ADO.NET database server.



In this figure, the ODBC application includes code to export data to the CSV file, and the ADO.NET application includes code to specify and open the CSV file. Because the Control Center and the DataDirect ODBC drivers use a consistent format, interoperability is supported via these standard interfaces.

Zen Common Assembly

The Zen Bulk Load implementation for ADO.NET uses the de facto standard defined by the Microsoft SqlBulkCopy classes, and adds powerful built-in features to enhance interoperability as well as the flexibility to make bulk operations more reliable.

The data provider includes provider-specific classes to support Zen Bulk Load. See [Data Provider-Specific Classes](#) for more information. If you use the Common Programming Model, you can use the classes in the Zen Common Assembly (see [Zen Common Assembly Classes](#)).

The Pervasive.Data.Common assembly includes classes that support Zen Bulk Load, such as the CsvDataReader and CsvDataWriter classes that provide functionality between bulk data formats.

The Common assembly also extends support for bulk load classes that use the Common Programming Model. This means that the SqlBulkCopy patterns can now be used in a new DbBulkCopy hierarchy.

Future versions of the data provider will include other features that enhance the Common Programming Model experience. See [Zen Common Assembly Classes](#) for more information on the classes supported by the Pervasive.Data.Common assembly.

Bulk Load Data File

The results of queries between dissimilar data stores are persisted in a comma-separated value (CSV) format file, a bulk load data file. The file name, which is defined by the BulkFile property, is using for writing and reading the bulk data. If the file name does not contain an extension, the ".csv" extension is assumed.

Example

The Zen source table GBMAXTABLE contains four columns. The following C# code fragment writes the GBMAXTABLE.csv and GBMAXTABLE.xml files that will be created by the CsvDataWriter. Note that this example uses the DbDataReader class.

```
cmd.CommandText = "SELECT * FROM GBMAXTABLE ORDER BY INTEGERCOL";
DbDataReader reader = cmd.ExecuteReader();
CsvDataWriter csvWriter = new CsvDataWriter();
csvWriter.WriteToFile(@"\NC1\net\Zen\GBMAXTABLE\GBMAXTABLE.csv", reader);
```

The bulk load data file GBMAXTABLE.csv contains the results of the query:

```
1,0x6263,"bc","bc"
2,0x636465,"cde","cde"
3,0x64656667,"defg","defg"
4,0x6566676869,"efghi","efghi"
5,0x666768696a6b,"fghijk","fghijk"
6,0x6768696a6b6c6d,"ghijklm","ghijklm"
7,0x68696a6b6c6d6e6f,"hijklmno","hijklmno"
8,0x696a6b6c6d6e6f7071,"ijklmnopq","ijklmnopq"
9,0x6a6b6c6d6e6f70717273,"jklmnopqrs","jklmnopqrs"
10,0x6b,"k","k"
```

The GBMAXTABLE.xml file, which is the bulk load configuration file that provides the format of this bulk load data file, is described in the following section.

Bulk Load Configuration File

A bulk load configuration file is produced when the CsvDataWriter.WriteToFile method is called (see [CsvDataWriter](#) for more information).

The bulk load configuration file defines the names and data types of the columns in the bulk load data file. These names and data types are defined the same way as the table or result set from which the data was exported.

If the bulk data file cannot be created or does not comply with the schema described in the XML configuration file, an exception is thrown. See [XML Schema Definition for a Bulk Data Configuration File](#) for more information about using an XML schema definition.

If a bulk load data file that does not have a configuration file is read, the following defaults are assumed:

-
- All data is read in as character data. Each value between commas is read as character data.
 - The default character set is the character set of the platform on which the Bulk Load CSV file is being read. See [Character Set Conversions](#) for more information.

The bulk load configuration file describes the bulk data file and is supported by an underlying XML schema.

Example

The format of the bulk load data file shown in the previous section is defined by the bulk load configuration file, GBMAXTABLE.xml. The file describes the data type and other information about each of the four columns in the table.

```
<?xml version="1.0" encoding="utf-8"?>
<table codepage="UTF-16LE" xsi:noNamespaceSchemaLocation="http://www.datadirect.com/ns/bulk/
BulkData.xsd" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <row>
    <column datatype="DECIMAL" precision="38" scale="0" nullable=
      "false">INTEGERCOL</column>
    <column datatype="VARBINARY" length="10" nullable=
      "true">VARBINCOL</column>
    <column datatype="VARCHAR" length="10" sourcecodepage="Windows-1252"
      externalfilecodepage="Windows-1252" nullable="true">VCHARCOL</column>
    <column datatype="VARCHAR" length="10" sourcecodepage="Windows-1252"
      externalfilecodepage="Windows-1252" nullable="true">UNIVCHARCOL</column>
  </row>
</table>
```

Determining the Bulk Load Protocol

Bulk operations can be performed using dedicated bulk protocol, that is, the data provider uses the protocol of the underlying database. In some cases, the dedicated bulk protocol is not available, for example, when the data to be loaded is in a data type not supported by the dedicated bulk protocol. Then, the data provider automatically uses a non-bulk method such as array binding to perform the bulk operation, maintaining optimal application uptime.

Character Set Conversions

At times, you might need to bulk load data between databases that use different character sets.

For the Control Center, the default source character data, that is, the output from the CsvDataReader and the input to the CsvDataWriter, is in Unicode (UTF-16) format. The source character data is always transliterated to the code page of the CSV file. If the threshold is exceeded and data is written to the external overflow file, the source character data is transliterated to the code page specified by the externalfilecodepage attribute defined in the bulk configuration XML schema (see [XML Schema Definition for a Bulk Data Configuration File](#)). If

the configuration file does not define a value for `externalfilecodepage`, the CSV file code page is used.

To avoid unnecessary transliteration, it's best for the CSV and external file character data to be stored in Unicode (UTF-16). You might want your applications to store the data in another code page in one of the following scenarios:

- The data will be written by ADO.NET and read by ODBC. In this case, the read (and associated transliteration) is done by ODBC. If the character data is already in the correct code page, no transliteration is necessary.
- Space is a consideration. Depending on the code page, the character data could be represented more compactly. For example, ASCII data is a single byte per character, UTF-16 is 2 bytes per character).

The configuration file may optionally define a second code page for each character column. When character data exceeds the value defined by the `CharacterThreshold` property and is stored in a separate file (see [External Overflow File](#)), the value defines the code page for that file.

If the value is omitted or if the code page defined by the source column is unknown, the code page defined for the CSV file will be used.

External Overflow File

If the value of the `BinaryThreshold` or `CharacterThreshold` property of the `CsvDataWriter` object is exceeded, separate files are generated to store the binary or character data. These overflow files are located in the same directory as the bulk data file.

If the overflow file contains character data, the character set of the file is governed by the character set specified in the CSV bulk configuration file.

The filename contains the CSV filename and a ".lob" extension (for example, `CSV_filename_nnnnnn.lob`). These files exist in the same location as the CSV file. Increments start at `_000001.lob`.

Bulk Copy Operations and Transactions

By default, bulk copy operations are performed as isolated operations and are not part of a transaction. This means there is no opportunity for rolling the operation back if an error occurs.

Zen allows bulk copy operations to take place within an existing transaction. You can define the bulk copy operation to be part of a transaction that occurs in multiple steps. Using this approach

enables you to perform more than one bulk copy operation within the same transaction, and commit or roll back the entire transaction.

Refer to the Microsoft online help topic "Transaction and Bulk Copy Operations (ADO.NET)" for information about rolling back all or part of the bulk copy operation when an error occurs.

Using Diagnostic Features

The .NET Framework provides a Trace class that can help end users identify the problem without the program having to be recompiled.

The Control Center delivers additional diagnostic capability:

- Ability to trace method calls
- Performance Monitor hooks that let you monitor connection information for your application

Tracing Method Calls

Tracing capability can be enabled either through environment variables or the PsqlTrace class. The data provider traces the input arguments to all of its public method calls, as well as the outputs and returns from those methods (anything that a user could potentially call). Each call contains trace entries for entering and exiting the method.

During debugging, sensitive data can be read, even if it is stored as a private or internal variable and access is limited to the same assembly. To maintain security, trace logs show passwords as five asterisks (*****).

Using Environment Variables

Using environment variables to enable tracing means that you do not have to modify your application. If you change the value of an environment variable, you must restart the application for the new value to take effect.

The following table describes the environment variables that enable and control tracing.

Environment Variable	Description
PVSW_NET_Enable_Trace	If set to 1 or higher, enables tracing. If set to 0 (the default), tracing is disabled.
PVSW_NET_Recreate_Trace	If set to 1, recreates the trace file each time the application restarts. If set to 0 (the default), the trace file is appended.
PVSW_NET_Trace_File	Specifies the path and name of the trace file.

Notes

- Setting `PVSW_NET_Enable_Trace = 1` starts the tracing process. Therefore, you must define the property values for the trace file before enabling the trace. Once the trace processing starts, the values of the other environment variables cannot be changed.
- If tracing is enabled and no trace file is specified by either the connection string option or the environment variable, the data provider saves the results to a file named `PVSW_NETTrace.txt`.

Using Static Methods

Some users may find that using static methods on the data provider's `Trace` class to be a more convenient way to enable tracing. The following C# code fragment uses static methods on the `.NET Trace` object to create a `PsqlTrace` class with a trace file named `MyTrace.txt`. The values set override the values set in the environmental variables. All subsequent calls to the data provider will be traced to `MyTrace.txt`.

```
PsqlTrace.TraceFile="C:\\MyTrace.txt";  
PsqlTrace.RecreateTrace = 1;  
PsqlTrace.EnableTrace = 1;
```

The trace output has the following format:

```
<Correlation#> <Timestamp> <CurrentThreadName>  
  <Object Address> <ObjectName.MethodName> ENTER (or EXIT)  
    Argument #1 : <Argument#1 Value>  
    Argument #2 : <Argument#2 Value>  
    ...  
  RETURN: <Method ReturnValue> // This line only exists for EXIT
```

where:

Correlation# is a unique number that can be used to match up ENTER and EXIT entries for the same method call in an application.

Value is the hash code of an object appropriate to the individual function calls.

During debugging, sensitive data can be read, even if it is stored as private or internal variable and access is limited to the same assembly. To maintain security, trace logs show passwords as five asterisks (*****).

PerfMon Support

The Performance Monitor (PerfMon) and VS Performance Monitor (VSPerfMon) utilities allow you to record application parameters and review the results as a report or graph. You can also use Performance Monitor to identify the number and frequency of CLR exceptions in your

applications. In addition, you can fine-tune network load by analyzing the number of connections and connection pools being used.

The data provider installs a set of PerfMon counters that let you tune and debug applications that use the data provider. The counters are located in the Performance Monitor under the category name Zen ADO.NET data provider.

The following table describes the PerfMon counters that you can use to tune connections for your application.

Counter	Description
Current # of Connection Pools	Current number of pools associated with the process.
Current # of Pooled and Non-Pooled Connections	Current number of pooled and non-pooled connections.
Current # of Pooled Connections	Current number of connections in all pools associated with the process.
Peak # of Pooled Connections	The highest number of connections in all connection pools since the process started.
Total # of Failed Commands	The total number of command executions that have failed for any reason since the process started.
Total # of Failed Connects	The total number of attempts to open a connection that have failed for any reason since the process started.

For information on using PerfMon and performance counters, refer to the Microsoft documentation library.

Analyzing Performance With Connection Statistics

The .NET Framework 2.0 and higher supports run-time statistics, which are gathered on a per-connection basis. The Control Center supports a wide variety of run-time statistical items. These statistical items provide information that can help you to:

- Automate analysis of application performance
- Identify trends in application performance
- Detect connectivity incidents and send notifications
- Determine priorities for fixing data connectivity problems

Measuring the statistics items affects performance slightly. For best results, consider enabling statistics gathering only when you are analyzing network or performance behavior in your application.

Statistics gathering can be enabled on any Connection object, for as long as it is useful. For example, you can define your application to enable statistics before beginning a complex set of transactions related to performing a business analysis, and disable statistics when the task is complete. You can retrieve the length of time the data provider had to wait for the server and the number of rows that were returned as soon as the task is complete, or wait until a later time. Because the application disables statistics at the end of the task, the statistical items are measured only during the period in which you are interested.

Functionally, the statistical items can be grouped into four categories:

- Network layer items retrieve values associated with network activities, such as the number of bytes and packets that are sent and received and the length of time the data provider waited for replies from the server.
- Aggregate items return a calculated value, such as the number of bytes sent or received per round trip to the server.
- Row disposition statistical items provide information about the time and resources required to dispose of rows not read by the application.
- Statement cache statistical items return values that describe the activity of statements in a statement cache (see [Using Statement Caching](#) for more information on using the statement cache).

Enabling and Retrieving Statistical Items

When you create a Connection object, you can enable statistics gathering using the `StatisticsEnabled` property. The data provider begins the counts for the statistical items after you open a connection, and continues until the `ResetStatistics` method is called. If the connection is closed and reopened without calling `ResetStatistics`, the count on the statistical items continues from the point when the connection was closed.

Calling the `RetrieveStatistics` method retrieves the count of one or more statistical items. The values returned form a "snapshot in time" at the moment when the `RetrieveStatistics` method was called.

You can define the scope for the statistics gathering and retrieval. In the following C# code fragment, the statistical items measure only the Task A work; they are retrieved after processing the Task B work:

```
connection.StatisticsEnabled = true;
    // do Task A work
connection.StatisticsEnabled = false;
    // do Task B work
IDictionary currentStatistics = connection.RetrieveStatistics();
```

To view all the statistical items, you can use code like the following C# code fragment:

```
foreach (DictionaryEntry entry in currentStatistics) {
    Console.WriteLine(entry.Key.ToString() + ": " + entry.Value.ToString());
}
Console.WriteLine();
```

To view only the SocketReads and SocketWrites statistical items, you can use code like the following C# code fragment:

```
foreach (DictionaryEntry entry in currentStatistics) {
    Console.WriteLine("SocketReads = {0}",
        currentStatistics["SocketReads"]);
    Console.WriteLine("SocketWrites = {0}",
        currentStatistics["SocketWrites"]);
}
Console.WriteLine();
```

Note for Zen ADO.NET Entity Framework Users: The PsqlConnection methods and properties for statistics are not available at the Zen ADO.NET Entity Framework layer. Instead, the data provider exposes the same functionality through "pseudo" stored procedures. See [Using Stored Procedures with the ADO.NET Entity Framework](#) for more information.

The ADO.NET Data Providers

The Zen ADO.NET data providers deliver data access to any .NET-enabled application or application server. The Zen ADO.NET data providers enable point-to-point and *n*-tier access to data stores across the Internet and intranets. Because they are optimized for the .NET environment, the Zen ADO.NET data providers allow you to incorporate .NET technology and extend the functionality and performance of your existing system.

The following topics describe features that pertain to the Zen ADO.NET data providers:

- [About Zen ADO.NET Data Providers](#)
- [Using Connection Strings with the Zen ADO.NET Data Provider](#)
- [Performance Considerations](#)
- [Data Types](#)
- [Parameter Arrays](#)

Note: See [Zen ADO.NET Entity Framework Data Providers](#) for information about using the data provider with the Zen ADO.NET Entity Framework.

About Zen ADO.NET Data Providers

Zen ADO.NET data providers are built with 100% managed code, so they can run and connect to the database entirely within the common language runtime (CLR).

Code that runs in the native operating system, such as client libraries and COM components, is called unmanaged code. You can mix managed and unmanaged code within a single application. However, unmanaged code reaches outside the CLR, which means that it effectively raises complexity, reduces performance, and opens possible security risks.

Namespace

The namespace for the Zen ADO.NET data provider is `Pervasive.Data.SqlClient`. When connecting to the Zen database, you use the `PsqlConnection` and `PsqlCommand` objects in the `Pervasive.Data.SqlClient` namespace.

The following code fragment shows how to include a Zen ADO.NET data provider namespace in your applications:

C#

```
// Access Zen
using System.Data;
using System.Data.Common;
using Pervasive.Data.SqlClient;
```

Visual Basic

```
' Access Zen
Imports System.Data
Imports System.Data.Common
Imports Pervasive.Data.SqlClient
```

Assembly Name

The strongly named assembly for the Zen ADO.NET data provider is placed in the Global Assembly Cache (GAC) during installation. The assembly name is `Pervasive.Data.SqlClient.dll`.

The `Pervasive.Data.Common` assembly includes features such as support for bulk load.

Using Connection Strings with the Zen ADO.NET Data Provider

You can define the behavior of a connection using a connection string or the properties of the `PsqlConnection` object. However, values set in the connection string cannot be changed by the connection properties.

The basic format of a connection string includes a series of keyword/value pairs separated by semicolons. The following example shows the keywords and values for a simple connection string for the data provider:

```
"Server DSN=SERVERDEMO;Host=localhost"
```

See [Using Connection Strings](#) for guidelines on specifying connection strings.

Constructing a Connection String

`PsqlConnectionStringBuilder` property names are the same as the connection string option names. However, the connection string option name can consist of multiple words, with required spaces between the words. For example, the Min Pool Size connection string option is equivalent to the `MinPoolSize` property. [Using Connection Strings](#) lists the connection string properties, and describes each property.

The connection string options have the following form:

```
option name=value
```

Each connection string option value pair is separated by a semicolon. The following example shows the keywords and values for a simple connection string for the Zen ADO.NET data provider:

```
"Server DSN=SERVERDEMO;Host=localhost"
```

Performance Considerations

The performance of your application can be affected by the values you set for connection string options and the properties of some data provider objects.

Connection String Options that Affect Performance

Encrypt: Data encryption may adversely affect performance because of the additional overhead, mainly CPU usage, required to encrypt and decrypt data.

Max Statement Cache Size: Caching all of the prepared statements that an application uses might appear to offer the best performance. However, this approach may come at a cost of database server memory if you implement statement caching with connection pooling. In this case, each pooled connection has its own statement cache, which may contain all of the prepared statements used by the application. All of the cached prepared statements are also maintained in database server memory.

Pooling: If you enable the data provider to use connection pooling, you can define additional options that affect performance:

- **Load Balance Timeout:** You can define how long to keep connections in the pool. The pool manager checks a connection's creation time when it is returned to the pool. The creation time is compared to the current time, and if the timespan exceeds the value of the Load Balance Timeout option, the connection is destroyed. The Min Pool Size option can cause some connections to ignore this value.
- **Connection Reset:** Resetting a re-used connection to the initial configuration settings impacts performance negatively because the connection must issue additional commands to the server.
- **Max Pool Size:** Setting the maximum number of connections that the pool can contain too low might cause delays while waiting for a connection to become available. Setting the number too high wastes resources.
- **Min Pool Size:** A connection pool is created when the first connection with a unique connection string connects to the database. The pool is populated with connections up to the minimum pool size, if one has been specified. The connection pool retains this number of connections, even when some connections exceed their Load Balance Timeout value.

Schema Options: Returning some types of database metadata can affect performance. To optimize application performance, the data provider prevents the return of performance-expensive database metadata such as procedure definitions or view definitions. If your application needs this database metadata, you can specifically request its return.

To return more than one type of the omitted metadata, specify either a comma-separated list of the names, or the sum of the hexadecimal values of the column collections that you want to return. For example, to return procedure definitions and view definitions, specify one of the following:

- Schema Option=ShowProcedureDefinitions, ShowViewDefinitions
- Schema Options=0x60

Statement Cache Mode: In most cases, enabling statement caching results in improved performance. To enable the caching of prepared statements (Command instances), set this option to Auto. Use this setting if your application has marked prepared statements for implicit inclusion in the statement cache, or has marked some statements for implicit inclusion and others for explicit inclusion. If you want the statement cache to include only prepared statements that are marked Cache, 1) set the StatementCacheBehavior property of the Command object to Cache and 2) set this option to ExplicitOnly.

Properties that Affect Performance

StatementCacheBehavior: If your application reuses prepared statements multiple times over an application's lifetime, you can influence performance by using a statement cache. This property identifies how a prepared statement (a Command object instance) is handled during statement caching.

When set to Cache, the prepared statement is included in the statement cache.

When set to Implicit and the Statement Cache Mode connection string option is set to Auto, the prepared statement is included in the statement cache.

When set to DoNotCache, the prepared statement is excluded from the statement cache.

You can use connection statistics to determine the effect that caching specific statements has on performance (see [Analyzing Performance With Connection Statistics](#)).

Data Types

The following topics cover the data types supported by the Zen ADO.NET Data Provider:

- [Mapping Zen Data Types to .NET Framework Data Types](#)
- [Mapping System.Data.DbTypes to PsqldbTypes](#)
- [Mapping .NET Framework Types to PsqldbType](#)
- [Data Types Supported with Stream Objects](#)

Mapping Zen Data Types to .NET Framework Data Types

The following table lists the data types supported by the Zen ADO.NET data providers and how they are mapped to the .NET Framework types. You can use the table to infer the data types that will be used when a DataSet is filled using a DataAdapter.

This table also identifies the proper accessors for the data when a DataReader object is used directly.

- The Zen Data Type column refers to the native type name.
- The PsqldbType column refers to the ADO.NET Data Provider's type enumeration. Generally, there is a one to one mapping between the native type and the PsqldbType. The Zen NUMBER data type, which can be either a decimal or a double, is an exception to this rule.
- The .NET Framework Type column refers to the base data types available in the Framework.
- The .NET Framework Typed Accessor column refers to the method that must be used to access a column of this type when using a DataReader.

Mapping of Zen Data Types

Zen Data Type	PsqlDbType	.NET Framework Type	.NET Framework Typed Accessor
AUTOTIMESTAMP	Timestamp	DateTime	GetDateTime()
BFLOAT4	BFloat4	Single	GetSingle()
BFLOAT8	BFloat8	Double	GetDouble()
BIGIDENTITY	BigInt	Int64	GetInt64()
BIGINT	BigInt	Int64	GetDecimal()
BINARY	Binary	Byte[]	GetBytes()
BIT	Bit	Byte[]	GetBytes()
CHAR	Char	String Char[]	GetString() GetChars()
CURRENCY	Currency	Decimal	GetDecimal()
DATE	Date	DateTime	GetDateTime()
DATETIME	DateTime ¹	DateTime	GetDateTime()
DECIMAL	Decimal	Decimal	GetDecimal()
DOUBLE	Double	Double	GetDouble()
FLOAT	Float	Double	GetDouble()
IDENTITY	Identity	Int32	GetInt32()
INTEGER	Integer	Int32	GetInt32()
LONGVARBINARY	LongVarBinary	Byte[]	GetBytes()
LONGVARCHAR	LongVarChar	Byte[]	GetBytes()
MONEY	Money	Decimal	GetDecimal()
NCHAR	NChar	String Char[]	GetString() GetChars()
NLONGVARCHAR	NLongVarChar	String Char[]	GetString() GetChars()
NUMERIC	Decimal	Decimal	GetDecimal()
NUMERICSA	DecimalSA	Decimal	GetDecimal()

Zen Data Type	PsqlDbType	.NET Framework Type	.NET Framework Typed Accessor
NUMERICSTS	DecimalSTS	Decimal	GetDecimal()
NVARCHAR	NVarChar	String Char[]	GetString() GetChars()
REAL	Real	Single	GetSingle()
SMALLIDENTITY	SmallIdentity	Int16	GetInt16()
SMALLINT	SmallInt	Int16	GetInt16()
TIME	Time	Timespan ²	GetValue()
TIMESTAMP, TIMESTAMP2	Timestamp	DateTime	GetDateTime()
TINYINT	TinyInt	SByte	GetByte()
UBIGINT	UBigInt	UInt64	GetUInt64()
UNIQUE_IDENTIFIER	UniqueIdentifier ¹	String	GetString()
UIINTEGER	UIInteger	UInt32	GetUInt32()
USMALLINT	USmallInt	UInt16	GetUInt16()
UTINYINT	UTinyInt	Byte	GetByte()
VARCHAR	VarChar	String Char[]	GetString() GetChars()

1 Supported in Zen 9.5 and higher

2 Depends on the setting of the timetype connect option.

Mapping Parameter Data Types

The type of the parameter is specific to each data provider. Zen ADO.NET data providers must convert the parameter value to a native format before sending it to the server. The best way for an application to describe a parameter is to use the data provider-specific type enumeration.

In generic programming circumstances, the data provider-specific type may not be available. When no provider-specific DB type has been specified, the data type will be inferred from either the System.Data.DbType or from the .NET Framework type of the parameter's value.

Zen ADO.NET data providers use the following order when inferring the data type of a parameter:

-
- The data provider uses the provider-specific data type if it has been specified.
 - The data provider infers the data type from the `System.Data.DbType` if it has been specified, but the provider-specific data type has not been specified.
 - The data provider infers the data type from the .NET Framework type if neither the provider-specific data type nor the `System.Data.DbType` have been specified.

Mapping `System.Data.DbTypes` to `PsqlDbType`s

The following table shows how the data provider infers its types if only the `System.Data.DbType` is specified.

<code>System.Data.DbType</code>	<code>PsqlDbType</code>
<code>AnsiString</code>	<code>VarChar</code>
<code>AnsiStringFixedLength</code>	<code>Char</code>
<code>Binary</code>	<code>Binary</code>
<code>Boolean</code>	<code>Integer</code>
<code>Byte</code>	<code>Integer</code>
<code>Currency</code>	<code>Currency</code>
<code>Date</code>	<code>Date</code>
<code>DateTime</code>	<code>DateTime</code> ¹
<code>Decimal</code>	<code>Decimal</code> or <code>Money</code>
<code>Double</code>	<code>Double</code>
<code>Float</code>	<code>Float</code>
<code>GUID</code>	<code>UniqueIdentifier</code> *
<code>Int16</code>	<code>SmallInt</code>
<code>Int32</code>	<code>Integer</code>
<code>Int64</code>	<code>BigInt</code>
<code>Int64</code>	<code>BigIdentity</code>
<code>Sbyte</code>	<code>Integer</code>
<code>Single</code>	<code>BFloat4</code>
<code>String</code>	<code>NVarChar</code>

System.Data.DbType	PsqlDbType
StringFixedLength	NChar
Time	Time
UInt16	USmallInt
UInt32	UInteger
UInt64	UBigInt
VarNumeric	Decimal
1 Supported in PSQL 9.5 and higher.	

Mapping .NET Framework Types to PsqlDbType

The following table shows the mapping that the data provider uses to infer a data type if neither the provider-specific data type nor the System.Data.DbType are provided.

.NET Framework Type	PsqlDbType
Boolean	Integer
Byte	Integer
Byte[]	Binary
DateTime	Timestamp
Decimal	Decimal
Double	Double
Int16	SmallInt
Int32	Integer
Int64	BigInt
Single	BFloat4
String	NVarChar VarChar (if PvTranslate=Nothing)
UInt16	USmallInt
UInt32	UInteger
UInt64	UBigInt

Data Types Supported with Stream Objects

Zen ADO.NET data providers support the use of streams as inputs to long data parameters with the data types listed in the following table.

Provider Data Type	Stream Type Supported
LONGVARBINARY	Stream
LONGVARCHAR	TextReader

See [Using Streams as Input to Long Data Parameters](#) for a discussion of using streams.

Using Streams as Input to Long Data Parameters

Allowing the use of noncontiguous memory to represent a very large binary or text value, such as a video clip or a large document, improves performance, functionality, and scalability.

Stream objects used to read binary data are derived from the `System.IO.Stream` object and use the Framework data type of `byte[]`:

- `System.IO.BufferedStream`
- `System.IO.FileStream`
- `System.IO.MemoryStream`
- `System.Net.Sockets.NetworkStream`
- `System.Security.Cryptography.CryptoStream`

Stream objects used to read text data are derived from the `System.IO.TextReader` object and use the Framework data type of `string`:

- `System.IO.StreamReader`
- `System.IO.StringReader`

To enable the use of streams, you set the `Value` property of the `PsqlParameter` object to a specific instance of the stream (see [PsqlParameter Object](#)). When the command is executed, the data provider reads from the stream to extract the value.

The examples shipped with the data provider include a code example on inserting data into `LONGVARCHAR` and `LONGVARBINARY` columns using randomly generated data. The example also shows how to use streaming objects as inputs to `LONGVARCHAR` and `LONGVARBINARY` columns.

Parameter Markers

Parameter markers, including parameter markers for stored procedures, are specified in Zen ADO.NET data providers by using the "?" symbol in SQL statements.

```
UPDATE emp SET job = ?, sal = ? WHERE empno = ?
```

Because parameters are not named, the bindings must occur in the order of the parameters in the statement. This means that the calls to the Add() method on the PsqlParameterCollection object (adding the Parameter objects to the collection) must occur in the order of the "?"s in the command text.

Parameter Arrays

Parameter array binding is typically used with INSERT statements to speed up the time needed to fill a table. An application can specify rows of parameter values with a single execution of a command. The values can then be sent to the database server in a single round trip (depending on the native capabilities of the backend database).

Zen ADO.NET data providers support input parameter arrays for INSERT and UPDATE statements.

Zen ADO.NET Core Data Providers

Zen ADO.NET Core data providers support data access to any .NET-enabled application or application server. It delivers high-performance point-to-point and *n*-tier access to industry-leading data stores across the Internet and intranets. Because they are optimized for the .NET environment, Zen ADO.NET Core data providers allow you to incorporate .NET Core technology and extend the functionality and performance of your existing system.

The following topics cover features of Zen ADO.NET data providers:

- [About Zen ADO.NET Core Data Providers](#)
- [Creating an Application in Visual Studio Using Zen ADO.Net Core DLL](#)
- [Creating a UWP Application in Visual Studio Using Zen ADO.Net Core Data Provider](#)
- [ADO.NET Data Provider Features Missing in Zen ADO.NET Core Data Provider](#)

Note: See [Zen ADO.NET Entity Framework Core Data Providers](#) for information about using the data provider with the ADO.NET Entity Framework Core.

About Zen ADO.NET Core Data Providers

Zen ADO.NET Core data providers are supported with:

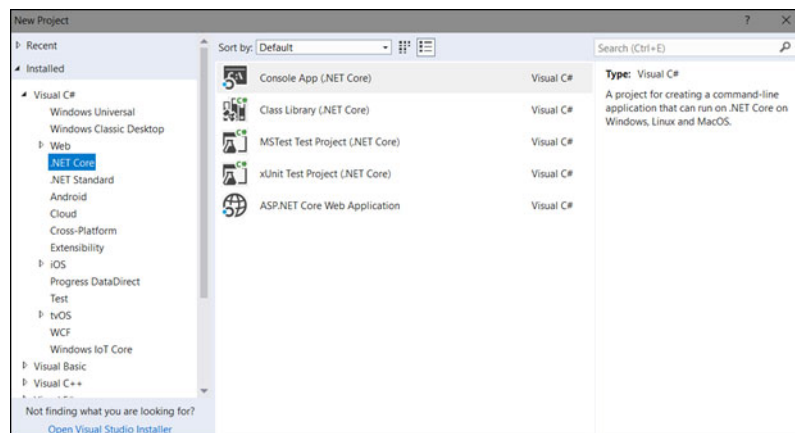
- Both the ADO.NET 4.4 and 4.5 providers support Visual Studio 2017 and 2019.
- The ADO.NET 4.5 provider also supports Visual Studio 2022.
- For the integration of the Zen data tools into Visual Studio, see the ADO.NET SDK readme.
- Visual Studio code
- .NET Standard 2.0, which supports .Net Core 2.0 +
For a full support matrix of .NET Standard 2.0, see <http://immo.landwerth.net/netstandard-versions/#>
- UWP applications

Creating an Application in Visual Studio Using Zen ADO.Net Core DLL

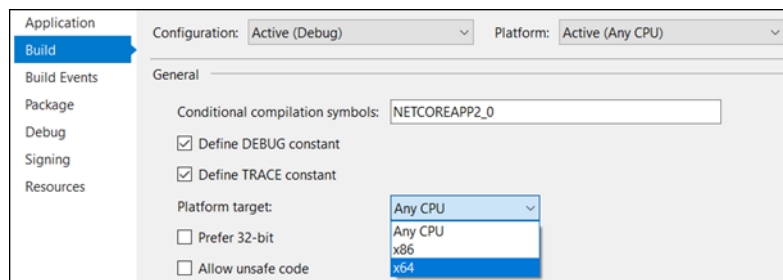
Before using these steps, first download the latest SDK .zip archive and extract the NuGet package Pervasive.Data.SqlClientStd for your version as listed under [Zen ADO.NET Data Providers Available with SDK Download](#).

To create an application in Visual Studio using Zen ADO.Net Core DLL

1. In Visual Studio, on the **File** menu, point to **New**, and then click **Project**. The New Project window appears.
2. In the left pane, in the **Installed** list, select the **.Net Core** template for Visual C#.



3. In the middle pane, select the required project type.
4. Enter the name and location for the project in the respective fields, then click **OK**.
5. Right-click the project, then click **Properties**.
6. In the **Build** pane, in the **Platform target list** select the required platform.



7. Add the downloaded NuGet package to the project.

To learn how to add the NuGet package locally, see
<https://stackoverflow.com/questions/10240029/how-do-i-install-a-nuget-package-nupkg-file-locally/38663739#38663739>

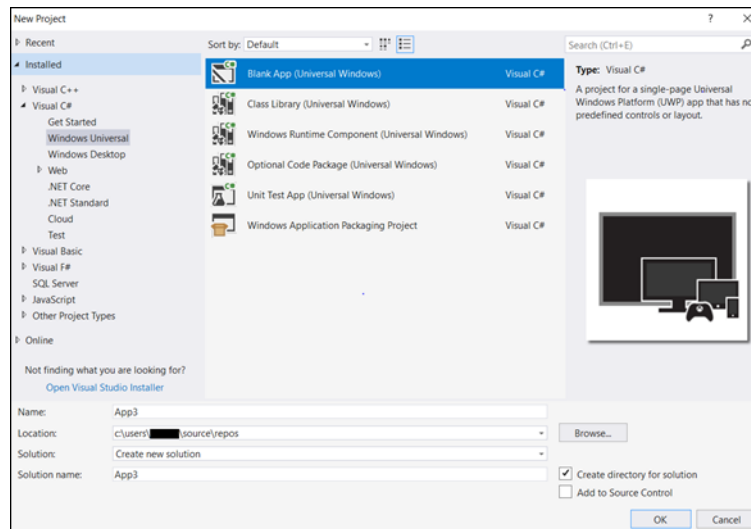
Creating a UWP Application in Visual Studio Using Zen ADO.Net Core Data Provider

Before using these steps, first download the latest SDK .zip archive and extract the NuGet package Pervasive.Data.SqlClientStd for your version as listed under [Zen ADO.NET Data Providers Available with SDK Download](#).


Note: For a UWP application to use Zen ADO.Net Core DLL, the required Windows 10 operating system version is 1709 (OS Build 16299) or later.

To create a UWP application in Visual Studio using Zen ADO.Net Core data provider


1. In Visual Studio, on the **File** menu, point to **New**, and then click **Project**. The New Project window appears.
2. In the left pane, in the Installed list for Visual C#, select **Windows Universal**, then in the middle pane, select **Blank App (Universal Windows)**.



3. Enter the name and location for the project in the respective fields; then click **OK**.
4. In the Minimum Version list, select **Windows 10 Fall Creators Update (10.0; Build 16299)**.



5. In the MainPage.xaml file, add a button, and then rename the button to **Connect to Actian Zen**.



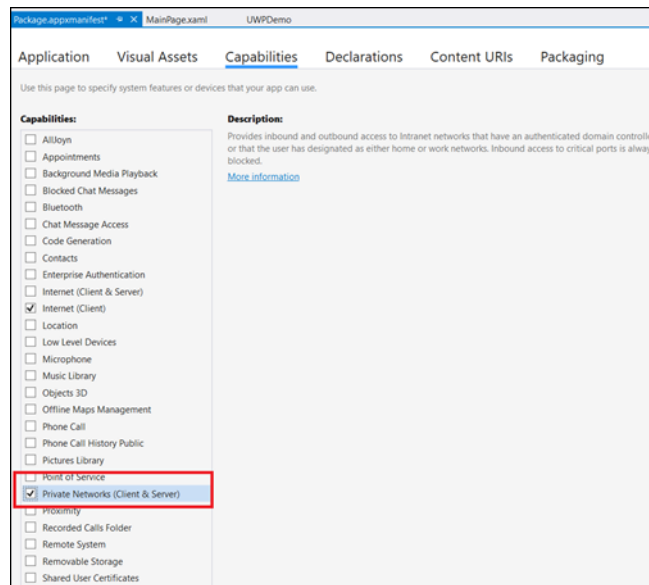
6. Double-click the button to open its implementation.
7. Add the following to the button implementation code:

```
private void Button_Click(object sender, RoutedEventArgs e)
{
    try
    {
        PsqlConnection conn = new PsqlConnection("Host=someHost;Port=1583;ServerDSN=DEFAULTDB;");
        conn.Open();
        Debug.WriteLine("Connection Opened: ");
    }
    catch (Exception eSql)
    {
        Debug.WriteLine("Exception Message: " + eSql.Message);
    }
}
```

8. Add the downloaded NuGet package to the project.

To learn how to add the NuGet package locally, see <https://stackoverflow.com/questions/10240029/how-do-i-install-a-nuget-package-nupkg-file-locally/38663739#38663739>

9. If the database is available in a remote machine, in the package.appxmanifest file on the Capabilities tab, select the **Private Networks (Client & Server)** check box.



10. Build and run the application.

ADO.NET Data Provider Features Missing in Zen ADO.NET Core Data Provider

The following ADO.NET data provider features are missing in Zen ADO.NET Core data providers:

- **Performance Counters:** As performance counters are not supported by ADO.Net Core, they also are not supported by Zen ADO.Net Core data providers.
- **Encoding:** As Zen ADO.Net data providers were created only for Windows-based applications, they use Windows-1252 encoding (identified by code page 1252) as the default encoding.

Since Zen ADO.Net Core data providers are a cross-platform, they use the default encoding of the current operating system as its default encoding.

- **CreatePermission in PervasiveFactory:** As CreatePermission (PermissionState) method is not supported by ADO.Net Core, it also is not supported by Zen ADO.Net Core data providers.
- **Fill Schema:** Zen ADO.Net data providers use Encoding.BodyName for Char, VarChar, and LongVarChar columns while pushing the data to the user. Since Encoding.BodyName is not supported by ADO.Net Core, Zen ADO.Net Core data providers use Encoding.WebName.

Zen ADO.NET Entity Framework Data Providers

Zen ADO.NET Entity Framework is an object-relational mapping (ORM) framework for the .NET Framework. Developers can use it to create data access applications by programming against a conceptual application model instead of directly against a relational storage schema. This model allows developers to decrease the amount of code to be written and maintained in data-centric applications.

Zen ADO.NET Entity Framework data providers (formerly Pervasive ADO.NET Entity Framework data providers) can be used with applications that use the ADO.NET Entity Framework.

Zen ADO.NET Entity Framework data providers are compatible with versions 6.1, 6.1.1, and 6.1.2 of the Microsoft ADO.NET Entity Framework. They support the following programming features:

- Applications targeting the .NET Framework versions listed under [Zen ADO.NET Data Providers Available with SDK Download](#).
- Database First, Code First, and Model First workflows
- Enumerated type support in all workflows
- Code First migrations
- "Plain-old" CLR objects (POCO) entities
- DbContext class
- Multiple DbContext classes
- Code First mapping to Insert, Update, and Delete stored procedures
- Configurable migration history
- Connection resiliency
- Index Attribute for Code First Migrations
- Disable Transactions for Function Imports
- Enum.HasFlag Support
- Allow Migrations commands to use context from reference instead of project
- Interceptors in web/app.config and DatabaseLogger

-
- Support for identifiers starting with '_'
 - Select concatenated string and numeric property

Zen ADO.NET Entity Framework data providers use ADO.NET data providers to communicate with the ADO.NET database server. This means that the functionality defined by Zen ADO.NET data providers applies to Zen ADO.NET Entity Framework data providers unless otherwise noted here. Similarly, any performance configurations made to Zen ADO.NET data providers are realized by Zen ADO.NET Entity Framework data providers.

Visual Studio 2017 or later is required when developing applications for the Zen ADO.NET Entity Framework. If you have configured Microsoft ADO.NET Entity Framework 6.1 (EF 6.1) to use Visual Studio 2017, you must install Entity Framework Tools 6.1.3 for Visual Studio 2017. However, once you install it, all your Visual Studio 2017 applications that previously used Microsoft ADO.NET Entity Framework 5.0 (EF5) must be upgraded to EF 6.1, after which you cannot revert to EF5.

About Zen ADO.NET Entity Framework Data Providers

Zen ADO.NET Entity Framework data providers are built with 100% managed code, so they can run and connect to the database entirely within the common language runtime (CLR).

Code that runs in the native operating system, such as client libraries and COM components, is called unmanaged code. You can mix managed and unmanaged code within a single application. However, unmanaged code reaches outside the CLR, which means that it effectively raises complexity, reduces performance, and opens possible security risks.

Namespace

The namespace for the Zen ADO.NET Entity Framework data providers is `Pervasive.Data.SqlClient.Entity`.

Note: The `Pervasive.Data.SqlClient.Entity` namespace is common to Microsoft ADO.NET Entity Framework versions 5.0 (EF 5) and 6.1 (EF 6.1).

Assembly Names

Zen ADO.NET Entity Framework data providers use the assembly name `Pervasive.Data.SqlClient.Entity.dll`.

To refer to EF 6.1 using the 4.4 version of the provider, select:

```
%windir%\Microsoft.NET\assembly\GAC_MSIL\Pervasive.Data.SqlClient.Entity\v4.0_4.4.0.6__c84cd5c63851e072
```

To refer to EF 6.1 using the 4.5 version of the provider, select:

```
%windir%\Microsoft.NET\assembly\GAC_MSIL\Pervasive.Data.SqlClient.Entity\v4.0_4.5.0.6__c84cd5c63851e072
```

Configuring Entity Framework 6.1

The Zen ADO.NET Entity Framework data provider supports the Microsoft ADO.NET Entity Framework versions 5.0 (EF5) and 6.1 (EF 6.1).

To use EF 6.1, you must first register it using one of the following methods:

- [Configuration File Registration](#)
- [Code-Based Registration](#)

Note: To register EF 6.1 while testing your applications locally, you can perform a code-based registration during development. However, when you deploy your project, you must perform a configuration file registration.

Configuration File Registration

To configure EF 6.1 by updating the configuration file

1. Install the **EntityFramework 6.1.2 NuGet** package.

An app.config file is created.

2. Remove the **defaultConnectionFactory** registration section from the app.config file and replace it with the following code according to the version of the provider you are using.

- For the 4.4 version:

```
<providers>
<provider invariantName="Pervasive.Data.SqlClient"
type="Pervasive.Data.SqlClient.Entity.PsqlProviderServices,
Pervasive.Data.SqlClient.Entity, Version=4.4.0.6, Culture=neutral,
PublicKeyToken=c84cd5c63851e072" />
</providers>
```

- For the 4.5 version:

```
<providers>
<provider invariantName="Pervasive.Data.SqlClient"
type="Pervasive.Data.SqlClient.Entity.PsqlProviderServices,
Pervasive.Data.SqlClient.Entity, Version=4.5.0.6, Culture=neutral,
PublicKeyToken=c84cd5c63851e072" />
</providers>
```

The EF 6.1 provider registration is added to Entity Framework section of the app.config file.

Code-Based Registration

To configure EF 6.1 through a code-based registration

1. Add the following new *DbConfiguration* class to your test application:

```
public class MyConfiguration : DbConfiguration
{
    public MyConfiguration()
    {
        SetProviderServices("PsqlProviderServices.ProviderInvariantName, new PsqlProviderServices());
    }
}
```

2. Add the following annotation on top of the *DbContext* class:

```
[DbConfigurationType(typeof(MyConfiguration))]
```

Using Multiple Entity Framework Versions Against the Same Database

A single database can use multiple versions of the Microsoft ADO.NET Entity Framework: 5.0 (EF5) and 6.1 (EF 6.1). However, when you switch between EF5 and EF6 applications against the same database, you will receive an error when you try saving to the database.

The error occurs due to the difference between the structure of the "__MigrationHistory" table used by EF5 and EF6.

To use EF5 and EF 6.1 applications against the same database without any errors, run the following command in the database:

```
drop table "__MigrationHistory"
```

Using Connection Strings with the Zen ADO.NET Entity Framework Data Provider

The Zen ADO.NET Entity Framework uses information contained in connection strings to connect to the underlying ADO.NET data provider that supports the Entity Framework. The connection strings also contain information about the required model and mapping files.

The data provider uses the connection string when accessing a model and mapping metadata and connecting to the data source.

You can specify an existing connection in the Entity Framework Wizard, or can define a new connection. Connection string options can be defined directly in a connection string, or set in the Advanced Properties dialog box in Visual Studio (see [Adding Connections in Server Explorer](#)).

Defining Connection String Values in Server Explorer

See [Adding Connections in Server Explorer](#) for detailed information about using Visual Studio to add and modify connections.

See Table for a description of the supported connection string options.

Changes in Default Values for Connection String Options

Most default values of the connection string options used by the ADO.NET Entity Framework data provider are the same as those used by the Zen ADO.NET data provider. The following table lists the connection string options that have a different default value when used with an ADO.NET Entity Framework application.

Connection String Option	Default Value in ADO.NET Entity Framework Application
Parameter Mode	Not supported.
Statement Cache Mode	ExplicitOnly is the only supported value.

Code First and Model First Support

Entity Framework 4.1 and later provide support for the Model First and Code First features. Implementing support for these features requires changes to the data provider, such as the way that long identifier names are handled. However, these changes do not require changes to your application.

Code First and Model First implementations require type mapping changes. See [Mapping Data Types and Functions](#) for more information.

Handling Long Identifier Names

Most Zen identifiers have a maximum length of 20 bytes. The identifier name can exceed this size because the names of the objects to be created on the server are taken from the class and property names. In addition, constraint names are often created by concatenating several object names. In these cases, the chances of exceeding the maximum identifier length are even greater.

The data provider shortens identifiers to database-allowed maximum identifier length, replacing the end of the identifier with an integer hash-code, for example, the string `ColumnMoreThanTwentyCharacters` is shortened to `ColumnMor_2873286151`. If you access or view the DB object using a DB tool, the names of the created tables may differ from what you might expect based on the Plain Old CLR Object (POCO) class names and property names (Code First), or the entity names and entity property names (Model First).

Note that when two identifiers that have the same leading characters are shortened, the difference between the identifiers is less obvious to a visual inspection. For example, assume that a table has two supporting sequences, `ColumnMoreThanTwentyCharacters` and `ColumnMoreThanTwenty1Characters`. When these sequences are shortened, they are renamed `ColumnMor_2873286151` and `ColumnMor_672399971`.

Using Code First Migrations with the ADO.NET Entity Framework

Entity Framework 4.3 and later support Code First Migrations, which enables you to update your database schema to reflect POCO classes without having to drop and recreate them.

Migrations enable you to incrementally evolve your database schema as your model changes. Each set of database changes is expressed in a code file, known as a migration. The migrations are ordered, typically using a time stamp, and a table in the database keeps track of which migrations are applied.

Code First Migrations implementation requires type mapping changes. See [Mapping Data Types and Functions](#) for more information.

To implement Code First Migrations using Progress DataDirect Connect for Zen ADO.NET data provider, you must perform the following additional settings:

1. Add references to the `Pervasive.Data.SqlClient.Entity` assembly in the project.

2. Inherit the Configuration Class changes and register the SQL Generator in the constructor of the Configuration Class. Do the following:

- Inherit the Configuration Class from `PervasiveDbMigrationsConfiguration <TContext>`. For example:

```
internal sealed class Configuration:  
PervasiveDbMigrationsConfiguration<%Context Name%>
```
- Register the Class Generator.

After you enable migrations using Package Manager Console, specify the Connection String either in the app.config or configuration.cs file along with additional settings in the configuration.cs file. However, if Connection String is specified in the app.config file, then ensure that the Connection String and the context have the same name.

If the Connection String is specified in the app.config file, use the following syntax to register SQL Generator in the app.config file according to the version of the provider you use:

- For the 4.4 version:

```
<providers>  
<provider invariantName="Pervasive.Data.SqlClient"  
type="Pervasive.Data.SqlClient.Entity.PsqlProviderServices,  
Pervasive.Data.SqlClient.Entity, Version=4.4.0.6, Culture=neutral,  
PublicKeyToken=c84cd5c63851e072" />  
</providers>
```

- For the 4.5 version:

```
<providers>  
<provider invariantName="Pervasive.Data.SqlClient"  
type="Pervasive.Data.SqlClient.Entity.PsqlProviderServices,  
Pervasive.Data.SqlClient.Entity, Version=4.5.0.6, Culture=neutral,  
PublicKeyToken=c84cd5c63851e072" />  
</providers>
```

To register SQL Generator in configuration.cs, use the following syntax:

```
SetSqlGenerator(PervasiveConnectionInfo.InvariantName, new PervasiveEntityMigrationSqlGenerator());
```

Using Enumerations with the ADO.NET Entity Framework

The `enum` keyword is used to declare an enumeration, a distinct type consisting of a set of named constants called the enumerator list. Every enumeration type has an underlying type. By default, every underlying type of the enumeration element is mapped to type `int32`. By default, the first enumerator has the value 0, and the value of each consecutive enumerator is incremented by 1. For example, you would specify a days-of-the-week enum type as:

```
enum Days {MON, TUE, WED, THU, FRI, SAT, SUN};
```

In this enumeration, `MON` would be 0, `TUE` 1, `WED` 2, and so forth. Enumerators can have initializers to override the default values. For example:

```
enum Days {MON=1, TUE, WED, THU, FRI, SAT, SUN};
```

In this enumeration, the sequence is forced to start at 1 instead of 0. The names of an enum type's fields are in uppercase letters., by convention, because they are constants.

Microsoft ADO.NET Entity Framework 5.0 and later support Enumerations. To use the enumeration feature, you must target .NET Framework 4.5 or later. Enumerations are supported in all three workflows, namely, Model First, Code First, and Database First.

In Entity Framework, an enumeration can have the following underlying types:

- `Byte`
- `Int16`
- `Int32`
- `Int64`
- `SByte`

By default, the enumeration is of type `Int32`. Another integral numeric type can be specified using a colon.

```
enum Days : byte{MON=1, TUE, WED, THU, FRI, SAT, SUN};
```

The underlying type specifies how much storage is allocated for each enumerator. However, an explicit cast is needed to convert from enum type to an integral type. Enum implementations also support type mapping changes. See [Mapping Data Types and Functions](#) for more information.

As part of Entity Framework, Entity Developer fully supports enum types by providing a new Enum node in its Model Explorer window. You can use the Enum property just like any other scalar property, such as in LINQ queries and updates.

Mapping Data Types and Functions

Developers can use the ADO.NET Entity Framework to create data access applications by programming against a conceptual application model instead of programming directly against a relational storage schema.

Type Mapping for Database First

In a Database First model, the data provider uses a store-centric type mapping scheme, in which the Zen (store) type influences the EDM type used when the model is generated.

Type Mapping for Database First

The following table shows Zen types are mapped to primitive types used in a Database First model. Some Zen data types can map to several different EDM types; the default values are shown in italics.

The columns are defined as follows:

- The Zen Type column refers to the native type name.
- The Store (SSDL) Type column refers to data types used by the store schema definition language (SSDL) file. The storage metadata schema is a formal description of the database that persists data for an application built on the EDM.
- The PrimitiveTypeKind column refers to the common data primitives used to specify the valid content of properties of entities used in defining EDM applications.

Zen Type	Store (SSDL) Type	PrimitiveTypeKind
AUTOTIMESTAMP	DateTime	DateTime
BFLOAT4	BFloat4	Single
BFLOAT8	BFloat8	Double
BIGIDENTITY	Bigint	Int64
BIGINT	Bigint	Int64
BINARY	binary	Byte[]
BIT	Bit	Boolean
CHAR	Char	String
CURRENCY	Currency	Decimal

Zen Type	Store (SSDL) Type	PrimitiveTypeKind
DATE	Date	DateTime
DATETIME	DateTime	DateTime
DECIMAL	Decimal	Decimal
DOUBLE	Double	Double
FLOAT	Float	Float
IDENTITY	Identity	Int32
INTEGER	Integer	Int32
LONGVARBINARY	LongVarBinary	Byte[]
LONGVARCHAR	LongVarChar	String
MONEY	Money	Decimal
NCHAR	NChar	String
NLONGVARCHAR	NLongVarChar	String
NUMERIC	Decimal	Decimal
NUMERICSA	DecimalSA	Decimal
NUMERICSTS	DecimalSTS	Decimal
NVARCHAR	NVarChar	String
REAL	Real	Single
ROWID	Rowid	Binary
SMALLIDENTITY	SmallIdentity	Int16
SMALLINT	Smallint	Int16
TIME	Time	Time
TIMESTAMP, TIMESTAMP2	DateTime	DateTime
TINYINT	TinyInt	SByte
UBIGINT	UBigInt	UInt64
UNIQUE_IDENTIFIER	Guid	Guid
UINTeger	UInteger	UInt32
USMALLINT	USmallInt	UInt16

Zen Type	Store (SSDL) Type	PrimitiveTypeKind
UTINYINT	UTinyInt	Byte
VARCHAR	Varchar	String

Type Mapping for Model First

The following table shows the model-centric type mapping where the EDM Simple Types influence the Zen (store) type used to create the database. The columns are defined as follows:

- The PrimitiveTypeKind column refers to the common data primitives used to specify the valid content of properties of entities used in defining EDM applications.
- Property Values Affecting Type Mapping identifies any property value that can affect type mapping.
- The Store (SSDL) column refers to data types used by the store schema definition language (SSDL) file. The storage metadata schema is a formal description of the database that persists data for an application built on the EDM.
- The Zen Type column refers to the native type name.

PrimitiveTypeKind	Property Values That Affect Type Mapping	Store (SSDL) Type	Zen Type
Binary	Fixed Length: TRUE	Binary	Binary(n)
	Fixed Length: FALSE	LongVarBinary	LongVarBinary
Boolean		Boolean	Bit
Byte		Tinyint_as_byte	TinyInt
DateTime		DateTime	DateTime
Decimal		Decimal	Decimal
Double		Double	Double
Guid		Guid	Guid
Single		Float	Float
SByte		Smallint_as_Sbyte	Smallint
Int16		SmallInt	Smallint
Int32		Integer	Integer
Int64		Bigint	BigInt

PrimitiveTypeKind	Property Values That Affect Type Mapping	Store (SSDL) Type	Zen Type
String	MaxLength= (1<=n<=8000) Fixed Length=True Unicode=False	Char	Char(n)
	MaxLength= (1<=n<=8000) Fixed Length=False Unicode=False	Varchar	Varchar(n)
	MaxLength= (>8000) Fixed Length=False Unicode=False	LongVarChar	LongVarchar
	MaxLength= (1<=n<=4000) Fixed Length=True Unicode=True	NChar	NChar(n)
	MaxLength= (1<=n<=4000) Fixed Length=False Unicode=True	NVarChar	NVarChar(n)
	MaxLength= (>4000) Fixed Length=False Unicode=True	NLongVarChar	NLongVarchar
Time		Time	Time
DateTimeOffset		DateTime	DateTime

Type Mapping for Code First

The following table shows the model-centric type mapping, where the CLR type influences the Zen (store) type used when the database is created. Some CLR types can map to several different Zen types; the default values are shown in *italics*.

The columns are defined as follows:

- The CLR Type column refers to the common language runtime type name.
- The Zen Type column refers to the native type name.

CLR Type	Zen Data Type
Byte[]	BINARY
Boolean	BIT
Byte	TINYINT
DateTime	DATETIME
Decimal	DECIMAL
Double	DOUBLE
Guid	UNIQUEIDENTIFIER BINARY
Single	FLOAT
Sbyte	SMALLINT
Int16	SMALLINT
Int32	INTEGER
Int64	BIGINT
String ¹	NCHAR NVARCHAR NLONGVARCHAR
TimeSpan	TIME
DateTimeOffset	DateTime

¹ In the Code First workflow, if the length of the string field in an entity is not specified, the data provider sets the default length to 2048 and 4096 bytes for unicode and non-unicode types respectively. However, if the length of the string field is set to a maximum allowed limit, that is 4000 bytes for unicode types and 8000 bytes for non-unicode types, the data provider resets it to 2048 bytes and 4096 bytes respectively. For all the other scenarios where the length of the string field is specified, the data provider uses the specified length.

Mapping EDM Canonical Functions to Zen Functions

The ADO.NET Entity Framework translates the Entity Data Model (EDM) canonical functions to the corresponding data source functionality for the ADO.NET Entity Framework data provider for Zen. The function invocations are expressed in a common form across data sources.

Because these canonical functions are independent of data sources, argument and return types of canonical functions are defined in terms of types in the EDM. When an Entity SQL query uses canonical functions, the appropriate function is called at the data source.

Both null-input behavior and error conditions are explicitly specified for all canonical functions. However, the ADO.NET Entity Framework does not enforce this behavior. Further details are available at: <http://msdn.microsoft.com/en-us/library/bb738626.aspx>

Aggregate Canonical Functions

The following table describes the mapping of EDM aggregate canonical functions to Zen functions.

Aggregate Canonical Function	Zen functions
<i>Avg(expression)</i>	<i>avg(expression)</i>
<i>BigCount(expression)</i>	<i>count(expression)</i>
<i>Count(expression)</i>	<i>count(expression)</i>
<i>Max(expression)</i>	<i>max(expression)</i>
<i>Min(expression)</i>	<i>min(expression)</i>
<i>StDev(expression)</i>	<i>stdev(expression)</i>
<i>StDevP(expression)</i>	<i>stdevp(expression)</i>
<i>Sum(expression)</i>	<i>sum(expression)</i>
<i>Var(expression)</i>	<i>var(expression)</i>
<i>VarP(expression)</i>	<i>varp(expression)</i>

Math Canonical Functions

The following table describes the mapping of EDM math canonical functions to Zen functions used to process columns that contain only decimal and integer values.

For more information, refer to the [Numeric Functions](#).

Math Canonical Function	Zen Function
<i>Abs(value)</i>	<i>abs(value)</i>
<i>Ceiling(value)</i>	<i>ceiling(value)</i>

Math Canonical Function	Zen Function
Floor(<i>value</i>)	floor(<i>value</i>)
Power(<i>value</i> , <i>exponent</i>)	power(<i>value</i> , <i>exponent</i>)
Round(<i>value</i>)	round(<i>numeric_expression1</i> , <i>integer_expression2</i>)
Round(<i>value</i> , <i>digits</i>)	round(<i>value</i> , <i>digits</i>)
Truncate(<i>value</i> , <i>digits</i>)	truncate(<i>value</i> , <i>digits</i>)

Date and Time Canonical Functions

The following table describes the mapping of EDM date and time canonical functions to Zen functions that generate, process, and manipulate data that consists of data types such as DATE and TIME.

Date and Time Canonical Function	Zen Functions
AddNanoseconds(<i>expression</i> , <i>number</i>)	dateadd(millisecond, <i>number</i> /1000000)
AddMicroseconds(<i>expression</i> , <i>number</i>)	dateadd(millisecond, <i>number</i> /1000)
AddMilliseconds(<i>expression</i> , <i>number</i>)	dateadd(millisecond, <i>number</i>)
AddSeconds(<i>expression</i> , <i>number</i>)	dateadd(second, <i>number</i>)
AddMinutes(<i>expression</i> , <i>number</i>)	dateadd(minute, <i>number</i>)
AddHours(<i>expression</i> , <i>number</i>)	dateadd(hour, <i>number</i>)
AddDays(<i>expression</i> , <i>number</i>)	dateadd(day, <i>number</i>)
AddMonths(<i>expression</i> , <i>number</i>)	dateadd(month, <i>number</i>)
AddYears(<i>expression</i> , <i>number</i>)	dateadd(year, <i>number</i>)
CreateDateTime(<i>year</i> , <i>month</i> , <i>day</i> , <i>hour</i> , <i>minute</i> , <i>second</i>)	datetimefromparts(<i>year</i> , <i>month</i> , <i>day</i> , <i>hour</i> , <i>minute</i> , <i>second</i> ,0)
CreateDateTimeOffset(<i>year</i> , <i>month</i> , <i>day</i> , <i>hour</i> , <i>minute</i> , <i>second</i> , <i>tzoffset</i>) ¹	datetimeoffsetfromparts(<i>year</i> , <i>month</i> , <i>day</i> , <i>hour</i> , <i>minute</i> , <i>second</i> , <i>tzoffset</i>)
CreateTime(<i>hour</i> , <i>minute</i> , <i>second</i>) ¹	timefromparts(<i>hour</i> , <i>minute</i> , <i>second</i> ,0,0)
CurrentDateTime()	now()
CurrentDateTimeOffset()	sysdatetimeoffset()
CurrentUtcDateTime()	current_timestamp()

Date and Time Canonical Function	Zen Functions
Day(expression)	datepart(day,expression)
DayOfYear(startexpression,endexpression)	dayofyear(expression)
DiffNanoSeconds(startexpression,endexpression)	datediff(millisecond,startexpression,endexpression) *1000000
DiffMilliSeconds(startexpression,endexpression)	datediff(millisecond,startexpression,endexpression)
DiffMicroSeconds(startexpression,endexpression)	datediff(millisecond,startexpression,endexpression) *1000
DiffSeconds(startexpression,endexpression)	datediff(second,startexpression,endexpression)
DiffMinutes(startexpression,endexpression)	datediff(minute,startexpression,endexpression)
DiffHours(startexpression,endexpression)	datediff(hour, startexpression,endexpression)
DiffDays(startexpression,endexpression)	datediff(day, startexpression, endexpression)
DiffMonths(startexpression,endexpression)	datediff(month,startexpression,endexpression)
DiffYears(startexpression,endexpression)	datediff(year,startexpression,endexpression)
GetTotalOffsetMinutes(DateTime Offset)	datepart(tzoffset,expression)
Year(expression)	datepart(year,expression)
Month(expression)	datepart(month,expression)
Day(expression)	datepart(day,expression)
Hour(expression)	datepart(hour,expression)
Minute(expression)	datepart(minute,expression)
Second(expression)	datepart(second,expression)
Millisecond(expression)	datepart(millisecond,expression)
TruncateTime(expression)	convert(expression, SQL_DATE)
¹ Requires Zen v11.30 Update 4 (May 2013)	

Bitwise Canonical Functions

The following table describes the mapping of EDM bitwise canonical functions to Zen functions.

Bitwise Canonical Function	Zen Functions
BitWiseAnd (value1, value2)	bit_and (value1, value2)

Bitwise Canonical Function	Zen Functions
BitWiseNot (<i>value</i>)	bit_compliment
BitWiseOr (<i>value1</i> , <i>value2</i>)	bit_or
BitWiseXor (<i>value1</i> , <i>value2</i>)	bit_xor

String Canonical Functions

The following table describes the mapping of EDM string canonical functions to Zen functions.

String Canonical Function	Zen Function
Concat(<i>string1</i> , <i>string2</i>)	concat(<i>string1</i> , <i>string2</i>)
Contains(<i>string</i> , <i>target</i>)	contains(<i>string</i> , <i>target</i>)
EndsWith(<i>string</i> , <i>target</i>)	endswith(<i>string</i> , <i>target</i>)
IndexOf(<i>target</i> , <i>string2</i>)	instr(<i>target</i> , <i>string2</i>)
Left(<i>string1</i> , <i>length</i>)	left(<i>string1</i> , <i>length</i>)
Length(<i>string</i>)	length(<i>string</i>)
LTrim(<i>string</i>)	ltrim(<i>string</i>)
Trim(<i>string</i>)	trim (BOTH FROM <i>string</i>)
Replace(<i>string1</i> , <i>string2</i> , <i>string3</i>)	replace(<i>string1</i> , <i>string2</i> , <i>string3</i>)
Reverse(<i>string</i>)	reverse(<i>string</i>)
RTrim(<i>string</i>)	rtrim(<i>string</i>)
StartsWith(<i>string</i> , <i>target</i>)	startswith(<i>string</i> , <i>target</i>)
Substring(<i>string</i> , <i>start</i> , <i>length</i>)	INCOMPLETE regexpr_substr(...)
ToLower(<i>string</i>)	lower(<i>string</i>)
ToUpper(<i>string</i>)	upper(<i>string</i>)

Other Canonical Functions

The following table describes the mapping of other canonical functions to Zen functions.

Other Canonical Function	Zen Function
NewGuid()	newid()

Extending Entity Framework Functionality

The ADO.NET Entity Framework offers powerful productivity gains by masking many ADO.NET features, simplifying application development. The ADO.NET Data Provider includes functionality designed to optimize performance.

Applications that use the standard Logging Application Block (LAB) from the Microsoft Enterprise Library 6.0 and the related design patterns can quickly display the SQL generated as part of the ADO.NET Entity Framework data providers.

See [Logging Application Blocks](#) for more information.

Enhancing Entity Framework Performance

Although the Entity Framework offers powerful productivity gains, some developers believe that the Entity Framework takes too much control of the features they need to optimize performance in their applications.

Limiting the Size of XML Schema Files

Building large models with the Entity Data Model (EDM) can be very inefficient. For optimal results, consider breaking up a model when it has reached 50 to 100 entities.

The size of the XML schema files is to some extent proportional to the number of tables, views, or stored procedures in the database from which you generated the model. As the size of the schema files increase, additional time is needed to parse and create an in-memory model for the metadata. This is a one-time performance cost that is incurred for eachObjectContext instance.

This metadata is cached per application domain, based on the EntityConnection String. This means that if you use the same EntityConnection string in multipleObjectContext instances in a single application domain, the application incurs the cost of loading metadata only once. However, the performance cost could still be significant if the size of the model becomes large and the application is not a long-running one.

Using Stored Procedures with the ADO.NET Entity Framework

Using stored procedures with the ADO.NET Entity Framework requires mapping functions. Calling these stored procedures is complex and requires some coding.

Providing Functionality

The Connection object includes properties and methods that provide enhanced statistics functionality that are standard in the ADO.NET data provider, but are not available at the ADO.NET Entity Framework layer. Instead, you expose the same functionality through "pseudo" stored procedures.

This approach uses the Entity Data Model (EDM) to achieve results that correspond to the ADO.NET results. This in effect provides entities and functions backed by pseudo stored procedures.

Mapping to Pseudo Stored Procedure

The following table lists the mapping of the data provider's Connection properties to the corresponding pseudo stored procedure.

Connection Property	Pseudo Stored Procedure
StatisticsEnabled	Psql_Connection_EnableStatistics Psql_Connection_DisableStatistics
Connection Method	Pseudo Stored Procedure
ResetStatistics	Psql_Connection_ResetStatistics
RetrieveStatistics	Psql_Connection_RetrieveStatistics

Applications must use theObjectContext to create a stored procedure command as shown in the following C# code fragment:

```
using (MyContext context = new MyContext())
{
    EntityConnection entityConnection = (EntityConnection)context.Connection;

    // The EntityConnection exposes the underlying store connection
    DbConnection storeConnection = entityConnection.StoreConnection;
    DbCommand command = storeConnection.CreateCommand();
    command.CommandText = "Psql_Connection_EnableStatistics";
    command.CommandType = CommandType.StoredProcedure;
    command.Parameters.Add(new SqlParameter("cid", 1));
}
```

```
//  
bool openingConnection = command.Connection.State == ConnectionState.Closed;  
if (openingConnection) { command.Connection.Open(); }  
int result;  
try  
{  
    result = command.ExecuteNonQuery();  
}  
finally  
{  
    if (openingConnection && command.Connection.State == ConnectionState.Open) {  
        command.Connection.Close(); }  
}
```

Using Overloaded Stored Procedures

If you have multiple overloaded stored procedures, the Zen Entity Framework data provider appends an identifier to each stored procedure name so you can distinguish between them in the SSDL. The data provider removes the appended identifier before calling the stored procedure for your application.

Using .NET Objects

The ADO.NET Entity Framework data provider supports the .NET public objects, exposing them as sealed objects.

For more information, see [.NET Objects Supported](#).

The ADO.NET Entity Framework programming contexts inherently eliminate the need to use some ADO.NET methods and properties. These properties and methods remain useful for standard ADO.NET applications. The online help, which is integrated into Visual Studio, describes the public methods and properties of each class.

Properties and Methods Differences with the ADO.NET Entity Data Provider

The following table lists the properties and methods that are not required or are implemented differently when using the data provider with an ADO.NET Entity application.

Property or Method	Behavior
PsqlCommand	
AddRowID	Not supported. The ADO.NET Entity Framework does not process the additional data that is returned.
ArrayBindCount	Not supported. The application cannot influence this bind count on top of the ADO.NET Entity Framework.
ArrayBindStatus	Not supported. The application cannot influence this bind count on top of the ADO.NET Entity Framework.
BindByName	Not supported. Instead, the data provider uses the ADO.NET Entity Framework programming contexts.
CommandTimeout	Not supported. Instead, the data provider uses the ADO.NET Entity Framework programming contexts.
UpdatedRowSource	Not supported. Instead, the data provider uses the ADO.NET Entity Framework programming contexts.
PsqlCommandBuilder	
DeriveParameters	Not supported. Instead, the data provider uses the ADO.NET Entity Framework programming contexts.

Property or Method	Behavior
PsqlConnection	
ConnectionTimeout	Supported only in a connection string.
StatisticsEnabled	Use the StatisticsEnabled or StatisticsDisabled stored procedure. See Using Stored Procedures with the ADO.NET Entity Framework for information on using this functionality in an ADO.NET Entity Framework application.
DataAdapter	
UpdateBatchSize	Not supported. Instead, the data provider uses the ADO.NET Entity Framework programming contexts.
Error	
ErrorPosition	Not supported. Instead, the data provider uses the ADO.NET Entity Framework programming contexts.
SQLState	Not supported. Instead, the data provider uses the ADO.NET Entity Framework programming contexts.

Creating a Model

The Entity Framework creates a model of your data in Visual Studio.

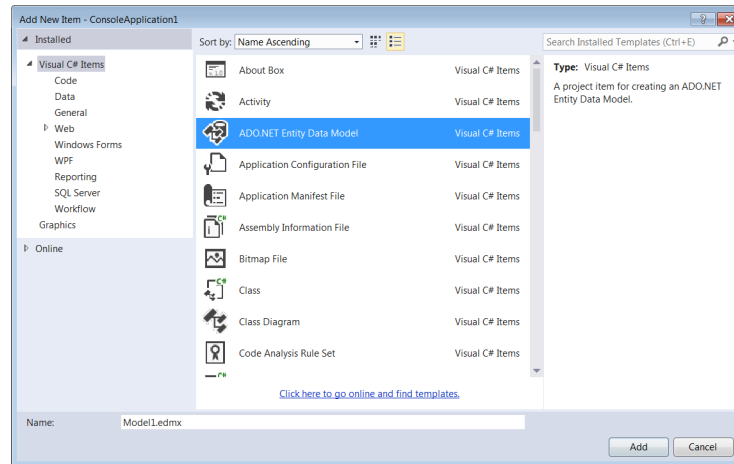
Note: Developing with the ADO.NET Entity Framework requires that you use Microsoft .NET Framework Version 4.5.x, 4.6.x, 4.7.x, or 4.8 and Visual Studio 2017 or later with the 4.4 or 4.5 version of the Actian Zen ADO.NET Entity Framework data provider.

To create a model of your data in Visual Studio using the Entity Framework, you must first ensure that you already have the database schema available.

To use the Entity Framework for creating a model of your data in Visual Studio


1. Create a new .NET application, such as Windows Console, Windows Forms, in Visual Studio.
2. In the Solution Explorer, right-click the project and select **Add > New Item**.

3. Select **ADO.NET Entity Data Model**, then click **Add**.

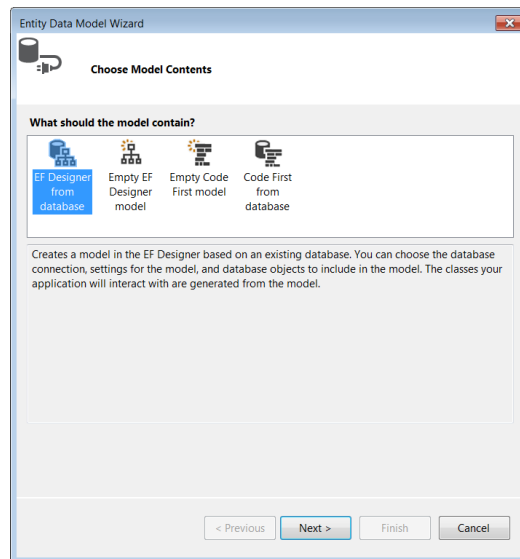


4. The Entity Data Model Wizard appears. Based on whether you have configured Microsoft ADO.NET Entity Framework 6.1 (EF 6.1), do one of the following:

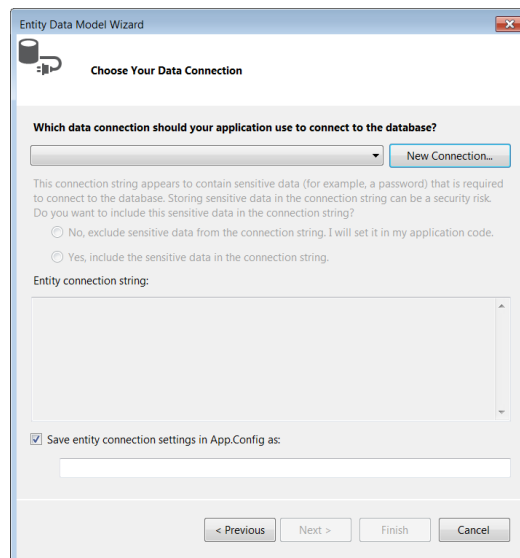
- If you **have not** configured EF 6.1, select **Generate from database** and click **Next**.



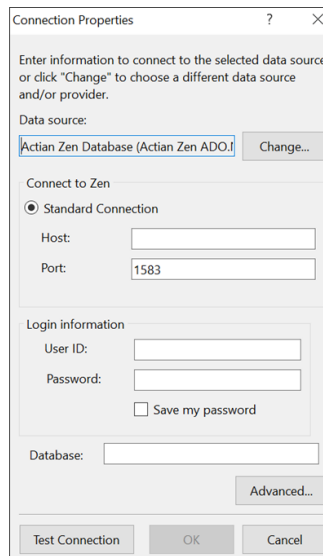
- If you **have** configured EF 6.1, select **EF Designer from database** and click **Next**.



5. On the Choose your Data Connection page, click **New Connection** to create a new connection. If you have an established connection, you can select it from the drop-down list.

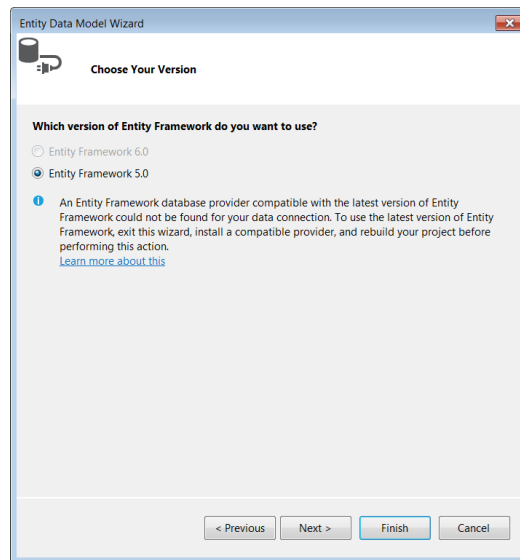


6. The Connection Properties window appears. Provide the necessary connection information and click **OK**.



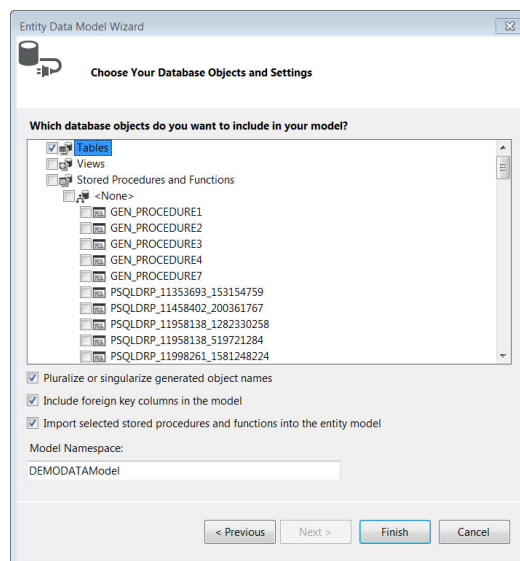
7. The Wizard creates an Entity connection string.
 - a. If the radio buttons are selectable, select **Yes, include the sensitive data in the connection string** to include the sensitive data in the connection string.
 - b. In the **Save entity connection settings** field, enter a name for the name of the main data access class or accept the default.
 - c. Click **Next**.
8. Based on the configured Entity Framework version, do one of the following:

- If you have configured EF5 for the current project, on the Choose Your Version page, proceed with the default **Entity Framework 5.0** by clicking **Next**.

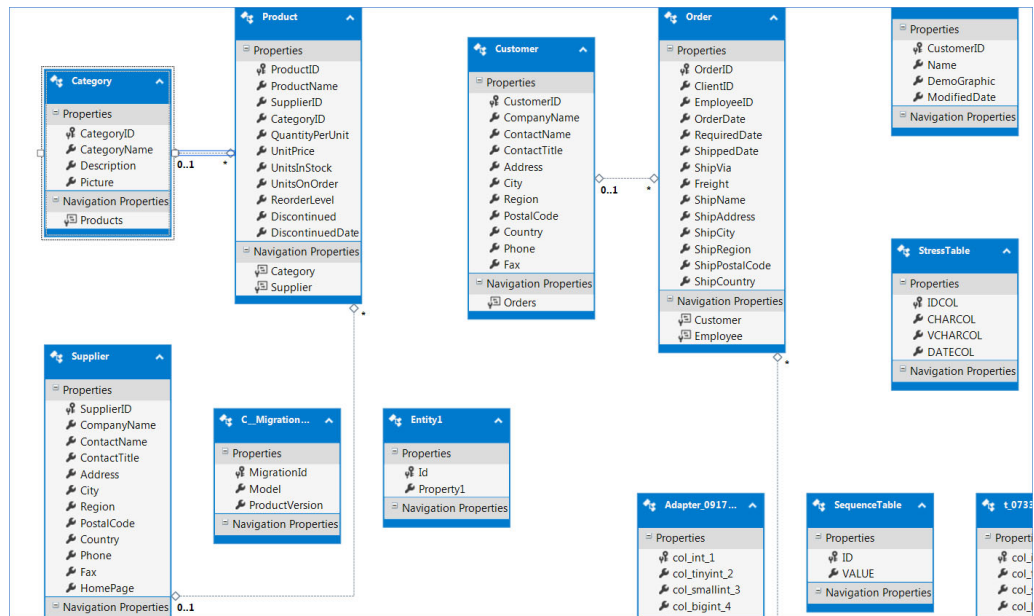


Note: To use the EF 6.1 with your current project, exit the wizard, configure EF 6.1, and then rebuild the project. When you rebuild the project after configuring EF 6.1, the wizard does not display the Choose Your Version page and you can directly proceed to the next step.

- If you have configured EF 6.1 for the current project, proceed to the next step.
9. Select the database objects that will be used in the model.



10. Click **Finish**. The model is generated and opened in the Model Browser.



Upgrading Entity Framework 5 Applications to Entity Framework 6.1

You can use Entity Framework Power Tools to regenerate Mapping Views with EF 6.1. For Entity Framework Power Tools to work with Zen ADO.NET Entity Framework Provider, you must add a Provider registration entry in the .NET Framework 4.0 machine.config file.

The steps given here upgrade an EF 5 application created using Zen ADO.NET Entity Framework provider to an EF 6.1 application. This procedure is needed only to upgrade an EF 5 application to EF 6.1. After upgrading, we recommend undoing machine.config file edits to restore their previous content.

Note: If the target database already contains objects created by the Entity Framework 5 code first application, then you need to drop these objects before running the migrated Entity Framework 6.1 application. Entity Framework 5.0 generates the foreign key constraint name differently from Entity Framework 6.1, which causes the application to fail with error "Table or view already exists."

To edit machine.config files

1. Close all Visual Studio windows before continuing with these steps.
2. Open the machine.config file for .Net Framework 4 in the following location, and also in the second one if you are using a 64-bit system:
 - %windir%\Microsoft.NET\Framework\v4.0.30319\Config\machine.config
 - %windir%\Microsoft.NET\Framework64\v4.0.30319\Config\machine.config
3. Under the <configSections></configSections> node, add the following entry:

```
<section name="entityFramework"
type="System.Data.Entity.Internal.ConfigFile.EntityFrameworkSection, EntityFramework,
Version=6.0.0.0, Culture=neutral, PublicKeyToken=b77a5c561934e089" requirePermission="false" />
```

4. Under the <configuration></configuration> node after the </configSections> closure tag, add the following entry:

```
<entityFramework>
<providers>
<provider invariantName="Pervasive.Data.SqlClient"
type="Pervasive.Data.SqlClient.Entity.PsqlProviderServices, Pervasive.Data.SqlClient.Entity,
Version=4.5.0.6, Culture=neutral, PublicKeyToken=c84cd5c63851e072" />
</providers>
</entityFramework>
```

Note: This version 4.5 example uses Version=4.5.0.6. For version 4.4, use Version=4.4.0.6.

5. Save and close the files you changed.
6. You must install EF version 6.1.2 to the corresponding project. To do so, you need to upgrade your project to the new EF 6.1.x runtime. You can do this by one of two methods:

Method 1

- a. Right-click your project and select **Manage NuGet Packages**.
- b. In the Online tab, select **EntityFramework** and click **Install**.

Any previous version of the EntityFramework NuGet package is upgraded to EF 6.1.x.

Method 2

Run the following command from Package Manager Console for the corresponding project to install EF 6.1.2 for that project:

```
Install-Package EntityFramework -Version 6.1.2
```

7. After using one of the methods, confirm removal of assembly references to System.Data.Entity.dll.

Installing the EF6 NuGet package should automatically remove any references to System.Data.Entity from your project.
8. Change any Enterprise Framework Designer (EDMX) models to use EF 6.x code generation. If you have any models created with the EF Designer, you will need to update the code generation templates to generate EF6-compatible code.

Note: For Visual Studio 2012 and later, only EF 6.x DbContext Generator templates are available.

- a. Delete existing code-generation templates:
These files are typically named <edmx_file_name>.tt and <edmx_file_name>.Context.tt and are nested under your .edmx file in Solution Explorer. You can select the templates in Solution Explorer and use the delete key to remove them.

Note: In website projects, the templates are not nested under your .edmx file, but alongside it in Solution Explorer. In VB.NET projects you need to enable Show All Files to see the nested template files.

- b. Add the appropriate EF 6.x code generation template:
Open your model in the EF Designer, right-click the design surface, and select Add Code Generation Item.
 - If you are using the DbContext API (recommended), then EF 6.x DbContext Generator is available under the Data tab.

Note that if you are using Visual Studio 2012, you need to install the EF 6 Tools to have this template. For details, see Get Entity Framework at Microsoft, [https://msdn.microsoft.com/en-us/library/ee712906\(v=vs.113\).aspx](https://msdn.microsoft.com/en-us/library/ee712906(v=vs.113).aspx).

- If you are using theObjectContext API, then you need to select the Online tab and search for EF 6.x EntityObject Generator.
 - c. If you applied any customizations to the code generation templates, you will need to reapply them to the updated templates.
9. Update name spaces for any core EF types you are using.

The name spaces for DbContext and Code First types have not changed. This means that for many applications that use EF 4.1 or later you will not need to change anything.

Types such as ObjectContext that were previously in System.Data.Entity.dll have been moved to new name spaces. This means you may need to update your using or import directives to build against EF6.

The general rule for namespace changes is that any type in System.Data.* is moved to System.Data.Entity.Core.*. In other words, just insert Entity.Core. after System.Data. For example:

- System.Data.EntityException => System.Data.Entity.Core.EntityException
- System.Data.Objects.ObjectContext => System.Data.Entity.Core.Objects.ObjectContext
- System.Data.Objects.DataClasses.RelationshipManager =>
System.Data.Entity.Core.Objects.DataClasses.RelationshipManager

These types are in the Core name spaces because they are not used directly for most DbContext-based applications. Some types that were part of System.Data.Entity.dll are still used commonly and directly for DbContext-based applications and so have not been moved into the Core name spaces. These are:

- System.Data.EntityState => System.Data.Entity.EntityState
- System.Data.Objects.DataClasses.EdmFunctionAttribute =>
System.Data.Entity.DbFunctionAttribute
Note: This class has been renamed. A class with the old name still exists and works but is now marked as obsolete.
- System.Data.Objects.EntityFunctions => System.Data.Entity.DbFunctions
Note: This class has been renamed. A class with the old name still exists and works but is now marked as obsolete.

10. Regenerate Mapping Views

If you have generated Mapping Views previously, delete the files and regenerate the Mapping View. For more information about Mapping Views, see <https://msdn.microsoft.com/en-us/data/dn469601>.

The EF 5 application built with the Zen ADO.NET data provider is now upgraded to EF 6.1.2.

Note: Once all of your EF 5 applications are upgraded to EF 6.1.2, we recommend that you undo the changes to Machine.config file and restore your earlier Machine.config file.

Entity Framework Power Tools provides an easier way to migrate EF 5.0 applications which use model Views to EF 6.1 applications. During our testing we found the Entity Framework Power Tools working fine with Zen Entity Framework Provider.

Note: Entity Framework Power Tools is not a certified or supported tool with Zen ADO.NET Entity Framework data providers.

For More Information

See the following sources for additional information about ADO.NET and the Entity Framework:

- *Programming Entity Framework* by Julie Lerman provides a comprehensive discussion of using the ADO.NET Entity Framework.
- [ADO.NET Entity Framework](#) introduces the Entity Framework and provides links to numerous detailed articles.
- [Connection Strings \(Entity Framework\)](#) describes how connection strings are used by the Entity Framework. The connection strings contain information used to connect to the underlying ADO.NET data provider as well as information about the required Entity Data Model mapping and metadata.
- [Entity Data Model Tools](#) describes the tools that help you to build applications graphically with the EDM: the Entity Data Model Wizard, the ADO.NET Entity Data Model Designer (Entity Designer), and the Update Model Wizard. These tools work together to help you generate, edit, and update an Entity Data Model.
- [LINQ to Entities](#) enables developers to write queries against the database from the same language used to build the business logic.

Zen ADO.NET Entity Framework Core Data Providers

Zen ADO.NET Entity Framework (EF) Core is a lightweight and extensible object-relational mapper (O/RM) for .NET that supports cross-platform development. It allows developers to work with the database using .NET objects and helps them decrease the amount of code that needs to be written and maintained in data-centric applications.

Zen ADO.NET Entity Framework Core data providers can be used with applications that use the ADO.NET Entity Framework Core. They support the following:

- Microsoft ADO.NET Entity Framework Core 3.1
- All the platforms supported by .Net Standard 2.1. For details, see <https://docs.microsoft.com/en-us/dotnet/standard/net-standard>.
- All the features supported by both Microsoft Entity Framework Core and Actian Zen database. For details, see <https://docs.microsoft.com/en-in/ef/core/what-is-new/>.

Zen ADO.NET Entity Framework Core data providers use ADO.NET data providers to communicate with ADO.NET database servers. This means that the functionality defined by Zen ADO.NET data providers applies to Zen ADO.NET Entity Framework Core data providers unless otherwise noted here. Similarly, any performance configurations made to Zen ADO.NET data providers are realized by the Zen ADO.NET Entity Framework Core data providers.

Note: Visual Studio 2019 or later is required when developing applications for the Zen ADO.NET Entity Framework Core.

About Zen ADO.NET Entity Framework Core Data Providers

Zen ADO.NET Entity Framework Core data providers are built with 100% managed code, so they can run and connect to the database entirely within the common language runtime (CLR).

Code that runs in the native operating system, such as client libraries and COM components, is called unmanaged code. You can mix managed and unmanaged code within a single application. However, unmanaged code reaches outside the CLR, which means that it effectively raises complexity, reduces performance, and opens possible security risks.

Namespace

The namespace for the Zen ADO.NET Entity Framework Core data provider is `Action.EntityFrameworkCore.Zen`.

Assembly Names

The Zen ADO.NET Entity Framework Core data provider uses the assembly name `Action.EntityFrameworkCore.Zen.dll`.

To use it, download the latest SDK .zip archive and extract the NuGet package `Action.EntityFrameworkCore.Zen` for your version as listed under [Zen ADO.NET Data Providers Available with SDK Download](#). Then add the package to your project.

Configuring the Zen ADO.NET Entity Framework Core Data Provider

To configure Zen ADO.NET Entity Framework Core Data Provider

1. Create an application that targets .Net Standard 2.1. For details, see <https://docs.microsoft.com/en-us/dotnet/standard/net-standard>.
2. Install the following NuGet package: **Action.EntityFrameworkCore.Zen**.
3. Add a new context class to your application and override the OnConfiguring method using the following code:

```
public class MyContext : DbContext
{
    protected override void OnConfiguring(DbContextOptionsBuilder optionsBuilder)
    => optionsBuilder.UseZen(connection string);
}
```

Using Connection Strings with the Zen ADO.NET Entity Framework Core Data Provider

The Zen ADO.NET Entity Framework Core data provider uses information contained in connection strings to connect to the underlying Zen ADO.NET data provider that supports the Entity Framework Core.

Changes in Default Values for Connection String Options

Most default values of the connection string options used by the Zen ADO.NET Entity Framework Core data provider are the same as those used by the Zen ADO.NET data provider (see the following table for more information). Table lists the connection string options that have a different default value when used with an ADO.NET Entity Framework Core application.

Connection String Option	Default Value in ADO.NET Entity Framework Core Application
Parameter Mode	Not supported.
Statement Cache Mode	ExplicitOnly is the only supported value.
DB File Directory Path	Environment.SpecialFolder.CommonApplicationData. To know more about the Environment.SpecialFolder enum and the supported fields, see https://docs.microsoft.com/en-us/dotnet/api/system.environment.specialfolder?view=netcore-3.1

Code First Support

Entity Framework Core supports Code First features. Implementing support for these features requires changes to the data provider, including those required to handle long identifier names. However, it does not need any changes to your application.

Code First implementations require type mapping changes. See [Type Mapping for Code First](#) for more information.

Handling Long Identifier Names

Most Zen identifiers have a maximum length of 20 bytes. The identifier name can exceed this size because the names of the objects to be created on the server are taken from the class and property names. In addition, constraint names are often created by concatenating several object names. In these cases, the chances of exceeding the maximum identifier length are even greater.

For columns, the data provider shortens identifiers to the maximum identifier length allowed by the database, replacing the end of the identifier with a tilde character ~. For example, the string `ColumnMoreThanTwentyCharacters` is shortened to `ColumnMoreThanTwent~`.

For tables, the data provider shortens identifiers to the maximum identifier length allowed by the database, replacing the end of the identifier with an integer hash-code. For example, the string `ATableWithAVeryLongTableName` is shortened to `ATableWit_1738385675`.

If you access or view the DB object using a DB tool, the names of the created tables may differ from what you might expect based on model class names and property names.

Using Code First Migrations with the ADO.NET Entity Framework Core

Entity Framework Core supports the Code First Migrations feature, which enables you to update your database schema to reflect model classes without having to drop and recreate them.

Migrations enable you to incrementally evolve your database schema as your model changes. Each set of changes to the database is expressed in a code file, known as a migration. The migrations are ordered, typically using a time stamp, and a table in the database keeps track of which migrations are applied to the database.

Code First Migrations requires type mapping changes. See [Type Mapping for Code First](#) for more information.

To implement Code First Migrations, once you have configured the Zen ADO.NET Entity Framework Core data provider, install the following NuGet package:

Microsoft.EntityFrameworkCore.Tools (version 3.1)

Using Reverse Engineering (Scaffolding)

The process of scaffolding entity type classes and a DbContext class based on a database schema is called reverse engineering. You can perform it using either the Scaffold-DbContext command of the EF Core Package Manager Console (PMC) tools or the `dotnet ef dbcontext scaffold` command of the .NET command-line interface (CLI) tools.

To use reverse engineering with the ADO.NET Action Zen Entity Framework Core data provider, perform the following steps after you have configured the data provider:

1. Install the following NuGet package: **Microsoft.EntityFrameworkCore.Tools (version 3.1)**.
2. Run the following Scaffold-DbContext PowerShell command:

```
Scaffold-DbContext 'connection string' Action.EntityFrameworkCore.Zen
```

You can add more parameters to the Scaffold-DbContext PowerShell command, if required. To know more about it, see <https://docs.microsoft.com/en-in/ef/core/managing-schemas/scaffolding?tabs=vs>.

To use reverse engineering in .Net Core CLI environment, see <https://docs.microsoft.com/en-in/ef/core/managing-schemas/scaffolding?tabs=dotnet-core-cli>.

Type Mapping for Code First

The following table shows the model-centric type mapping, where the CLR type influences the Zen (store) type used when the database is created. Some CLR types can map to several different Zen types.

The columns are defined as follows:

- The CLR Type column refers to the common language runtime type name.
- Property Values Affecting Type Mapping identifies any property value that can affect type mapping.
- The Zen Type column refers to the native type name.

CLR Type	Property Values That Affect Type Mapping	Action Zen Data Type
Bool		BIT
Byte		UTINYINT
Byte[]		BINARY
		LONGVARBINARY ¹
DateTime		AUTOTIMESTAMP
		DATE
		TIMESTAMP
		TIMESTAMP2 ²
DateTimeOffset		DATETIME
Decimal		CURRENCY
		DECIMAL ²
		NUMERIC
		NUMERICSA
		NUMERICSTS
		UBIGINT
Double		BFLOAT8
		DOUBLE ²

CLR Type	Property Values That Affect Type Mapping	Action Zen Data Type
Float		FLOAT ²
		REAL
		BFLOAT4
Guid		UNIQUEIDENTIFIER
Int		INTEGER ²
		IDENTITY
		USMALLINT
Long		BIGINT ²
		BIGIDENTITY
		UINTeger
SByte		TINYINT
Short		SMALLINT ²
		SMALLIDENTITY
String		NLONGVARCHAR ²
	Unicode=False	LONGVARCHAR
	MaxLength= (1<=n<=8000) Fixed Length=False Unicode=True	NVARCHAR
	MaxLength= (1<=n<=8000) Fixed Length=False Unicode=False	VARCHAR
	MaxLength= (1<=n<=8000) Fixed Length=True Unicode=True	NCHAR
	MaxLength= (1<=n<=8000) Fixed Length=True Unicode=False	CHAR
TimeSpan		TIME

CLR Type	Property Values That Affect Type Mapping	Action Zen Data Type
1. Maps to this type if no value is specified for MaxLength. 2. By default, maps to this type.		

Mapping EDM Canonical Functions to Zen Functions

The ADO.NET Entity Framework Core translates the Entity Data Model (EDM) canonical functions to the corresponding data source functionality for the Zen ADO.NET Entity Framework Core data provider. The function invocations are expressed in a common form across data sources.

Because these canonical functions are independent of data sources, argument and return types of canonical functions are defined in terms of types in the EDM. When an Entity SQL query uses canonical functions, the appropriate function is called at the data source.

Both null-input behavior and error conditions are explicitly specified for all canonical functions. However, the ADO.NET Entity Framework Core does not enforce this behavior.

Aggregate Canonical Functions

The following table describes the mapping of EDM aggregate canonical functions to Zen functions, and also the CLR types to which these functions apply.

Aggregate Canonical Function	Action Zen functions	CLR Type
BigCount(<i>expression</i>)	COUNT_BIG(<i>expression</i>)	Long
Count(<i>expression</i>)	COUNT(<i>expression</i>)	Int

Math Canonical Functions

The following table describes the mapping of EDM math canonical functions to Zen functions, and also the CLR types to which these functions apply.

Math Canonical Function	Action Zen Function	CLR Type
Abs(<i>expression</i>)	ABS(<i>expression</i>)	Decimal, Double, Float, Int, Long, SByte, Short
Ceiling(<i>expression</i>)	CEILING(<i>expression</i>)	Decimal, Double

Math Canonical Function	Action Zen Function	CLR Type
Floor(<i>expression</i>)	FLOOR(<i>expression</i>)	Decimal, Double
Pow(<i>base</i> , <i>power</i>)	POWER(<i>base</i> , <i>power</i>)	Double
Exp(<i>expression</i>)	EXP(<i>expression</i>)	Double
Log10(<i>expression</i>)	LOG10(<i>expression</i>)	Double
Log(<i>expression</i>)	LOG(<i>expression</i>)	Double
Sqrt(<i>expression</i>)	SQRT(<i>expression</i>)	Double
Acos(<i>expression</i>)	ACOS(<i>expression</i>)	Double
Asin(<i>expression</i>)	ASIN(<i>expression</i>)	Double
Atan(<i>expression</i>)	ATAN(<i>expression</i>)	Double
Atan2(<i>expression1</i> , <i>expression2</i>)	ATAN2(<i>expression1</i> , <i>expression2</i>)	Double
Cos(<i>expression</i>)	COS(<i>expression</i>)	Double
Sin(<i>expression</i>)	SIN(<i>expression</i>)	Double
Tan(<i>expression</i>)	TAN(<i>expression</i>)	Double
Sign(<i>expression</i>)	SIGN(<i>expression</i>)	Decimal, Double, Float, Int, Long, SByte, Short

Date and Time Canonical Functions

The following table describes the mapping of EDM date and time canonical functions to Zen functions that generate, process, and manipulate types that work with data time data, and also the CLR types to which these functions apply.

Date and Time Canonical Function	Action Zen Functions	CLR Type
<i>DateTime.Now</i>	<i>SYSDATETIME()</i>	DateTime
<i>DateTime.Now</i>	<i>SYSDATETIMEOFFSET ()</i>	DateTimeOffset
<i>DateTime.UtcNow</i>	<i>SYSUTCDATETIME()</i>	DateTime DateTimeOffset
<i>DateTime.Today</i>	<i>CURDATE()</i>	DateTime DateTimeOffset

Date and Time Canonical Function	Action Zen Functions	CLR Type
AddYears(<i>expression</i>)	DATEADD(<i>year</i> , <i>expression</i> , <i>column</i>)	DateTime DateTimeOffset
AddMonths(<i>expression</i>)	DATEADD(<i>month</i> , <i>expression</i> , <i>column</i>)	DateTime DateTimeOffset
AddDays(<i>expression</i>)	DATEADD(<i>day</i> , <i>expression</i> , <i>column</i>)	DateTime DateTimeOffset
AddHours(<i>expression</i>)	DATEADD(<i>hour</i> , <i>expression</i> , <i>column</i>)	DateTime DateTimeOffset
AddMinutes(<i>expression</i>)	DATEADD(<i>minute</i> , <i>expression</i> , <i>column</i>)	DateTime DateTimeOffset
AddSeconds(<i>expression</i>)	DATEADD(<i>second</i> , <i>expression</i> , <i>column</i>)	DateTime DateTimeOffset
AddMilliseconds(<i>expression</i>)	DATEADD(<i>millisecond</i> , <i>expression</i> , <i>column</i>)	DateTime DateTimeOffset
EF.Functions.DateDiffYear(<i>column</i> , <i>expression</i>)	DATEDIFF(<i>year</i> , <i>column</i> , <i>expression</i>)	DateTime DateTimeOffset
EF.Functions.DateDiffMonth(<i>column</i> , <i>expression</i>)	DATEDIFF(<i>month</i> , <i>column</i> , <i>expression</i>)	DateTime DateTimeOffset
EF.Functions.DateDiffDay(<i>column</i> , <i>expression</i>)	DATEDIFF(<i>day</i> , <i>column</i> , <i>expression</i>)	DateTime DateTimeOffset
EF.Functions.DateDiffHour(<i>column</i> , <i>expression</i>)	DATEDIFF(<i>hour</i> , <i>column</i> , <i>expression</i>)	DateTime DateTimeOffset TimeSpan
EF.Functions.DateDiffMinute(<i>column</i> , <i>expression</i>)	DATEDIFF(<i>minute</i> , <i>column</i> , <i>expression</i>)	DateTime DateTimeOffset TimeSpan
EF.Functions.DateDiffSecond(<i>column</i> , <i>expression</i>)	DATEDIFF(<i>second</i> , <i>column</i> , <i>expression</i>)	DateTime DateTimeOffset TimeSpan
EF.Functions.DateDiffMilliSecond(<i>column</i> , <i>expression</i>)	DATEDIFF(<i>millisecond</i> , <i>column</i> , <i>expression</i>)	DateTime DateTimeOffset TimeSpan

String Canonical Functions

The following table describes the mapping of EDM string canonical functions to Zen functions, and also the CLR types to which these functions apply.

String Canonical Function	Action Zen Functions	CLR Type
<code>IndexOf(<i>expression</i>)</code>	<code>POSITION(<i>expression</i>, <i>column</i>)</code>	String
<code>Replace(<i>toReplace</i>, <i>replaceWith</i>)</code>	<code>REPLACE(<i>toReplace</i>, <i>column</i>, <i>replaceWith</i>)</code>	String
<code>ToLower()</code>	<code>LOWER(<i>column</i>)</code>	String
<code>ToUpper()</code>	<code>UPPER(<i>column</i>)</code>	String
<code>SubString(<i>start</i>, <i>length</i>)</code>	<code>SUBSTRING(<i>column</i>, <i>start</i>, <i>length</i>)</code>	String
<code>IsNullOrWhiteSpace()</code>	A combination of LTRIM and RTRIM with a null check on the column	String
<code>TrimStart()</code>	<code>LTRIM(<i>column</i>)</code>	String
<code>TrimEnd()</code>	<code>RTRIM(<i>column</i>)</code>	String
<code>TRIM()</code>	A combination of LTRIM and RTRIM on the column	String
<code>Contains(<i>expression</i>)</code>	<code>POSITION(<i>expression</i>, <i>column</i>)</code>	String
<code>StartsWith(<i>expression</i>)</code>	A combination of LEFT and LENGTH on the column	String
<code>EndsWith(<i>expression</i>)</code>	A combination of RIGHT and LENGTH on the column	String
<code>Length()</code>	<code>LENGTH(<i>column</i>)</code>	String
<code>EF.Functions.Position(<i>column</i>, <i>expression</i>)</code>	<code>POSITION(<i>expression</i>, <i>column</i>)</code>	String

Other Canonical Functions

The following table describes the mapping of other canonical functions to Zen functions, and also the CLR types to which these functions apply.

Canonical Function	Action Zen Function	CLR Type
<code>ToString()</code>	<code>CONVERT(<i>column</i>, SQL_CHAR)</code>	All types
<code>NewGuid()</code>	<code>NEWID()</code>	Guid

Note: *Column* is the property to which the function is applied.

Extending Entity Framework Functionality

The ADO.NET Entity Framework Core and Actian Zen Entity Framework Core data provider are designed to be extended easily. The following examples demonstrate how to extend Entity Framework Core:

- <https://docs.microsoft.com/en-in/ef/core/managing-schemas/migrations/history-table>
- <https://docs.microsoft.com/en-in/ef/core/modeling/dynamic-model>

Using Stored Procedures with the ADO.NET Entity Framework Core

In Entity Framework Core, stored procedures can be executed using raw SQL queries. For more information, see <https://docs.microsoft.com/en-in/ef/core/querying/raw-sql>.

Upgrading an Application from Entity Framework 6.x to Entity Framework Core

To upgrade your application from Entity Framework 6.x to Entity Framework Core, see <https://docs.microsoft.com/en-us/ef/efcore-and-ef6/porting/>.

Limitations

The Actian Zen ADO.NET Entity Framework Core data provider has the following limitations:

- The Entity Framework Core has some limitations related to the reverse engineering (scaffolding) feature. They all apply to the Actian Zen Entity Framework Core data provider as well. For more information on these limitations, see <https://docs.microsoft.com/en-in/ef/core/managing-schemas/scaffolding?tabs=dotnet-core-cli>.
- The Actian Zen Entity Framework Core data provider does not support the concurrency feature of reverse engineering (scaffolding).

For More Information

Refer to the following sources for additional information about ADO.NET and the Entity Framework Core.

- [ADO.NET Entity Framework Core](#) introduces the Entity Framework Core and provides links to numerous detailed articles.
- [Feature Comparison](#) compares the features available in Entity Framework Core and Entity Framework 6.X.

[ASP.NET Core](#) shows how to use Entity Framework Core in an ASP.NET Core Razor Pages app.

Using Zen Data Providers in Visual Studio

Zen data providers support integration into Visual Studio. This means that developers can use the graphical user interface of Microsoft Visual Studio to perform a variety of tasks.

The following topics describe how the features of the Zen data providers are integrated into Visual Studio:

- [Adding Connections](#)
- [Using the Zen Performance Tuning Wizard](#)
- [Using Provider-Specific Templates](#)
- [Using the Zen Visual Studio Wizards](#)
- [Adding Components from the Toolbox](#)
- [Data Provider Integration Scenario](#)

Adding Connections

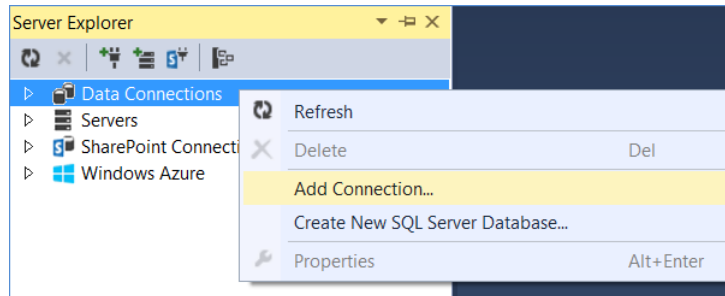
You can add connections in several ways in Visual Studio:

- [Adding Connections in Server Explorer](#)
- [Adding Connections with the Data Source Configuration Wizard](#)

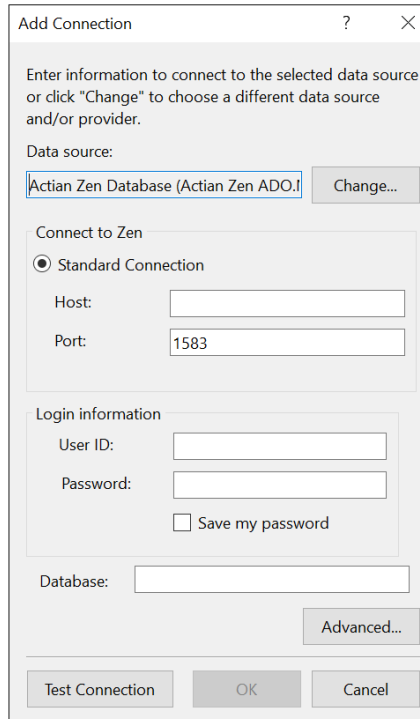
Adding Connections in Server Explorer

To add a connection

1. Right-click the **Data Connections** node in the Server Explorer and select **Add Connection**.

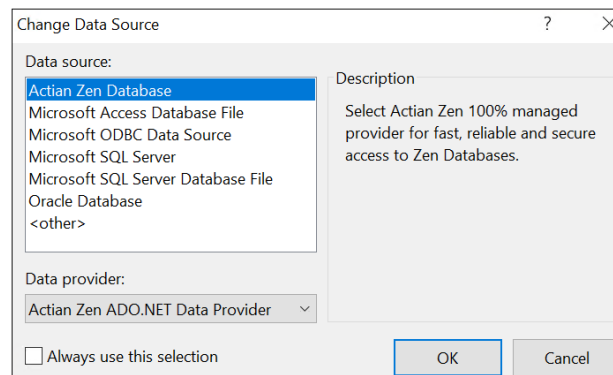


The Add Connection window appears.



The 'Add Connection' dialog box is shown. It has a title bar with a question mark and a close button. The main text says: 'Enter information to connect to the selected data source or click "Change" to choose a different data source and/or provider.' Below this is a 'Data source:' label and a text box containing 'Actian Zen Database (Actian Zen ADO.)'. To the right of the text box is a 'Change...' button. Below the 'Data source' section is a 'Connect to Zen' section with a radio button labeled 'Standard Connection' which is selected. Below this are 'Host:' and 'Port:' labels with corresponding text boxes. The 'Port' box contains '1583'. Below the 'Connect to Zen' section is a 'Login information' section with 'User ID:' and 'Password:' labels and corresponding text boxes. There is a checkbox labeled 'Save my password' which is unchecked. Below the 'Login information' section is a 'Database:' label and a text box. To the right of the 'Database' text box is an 'Advanced...' button. At the bottom of the dialog are three buttons: 'Test Connection', 'OK', and 'Cancel'.

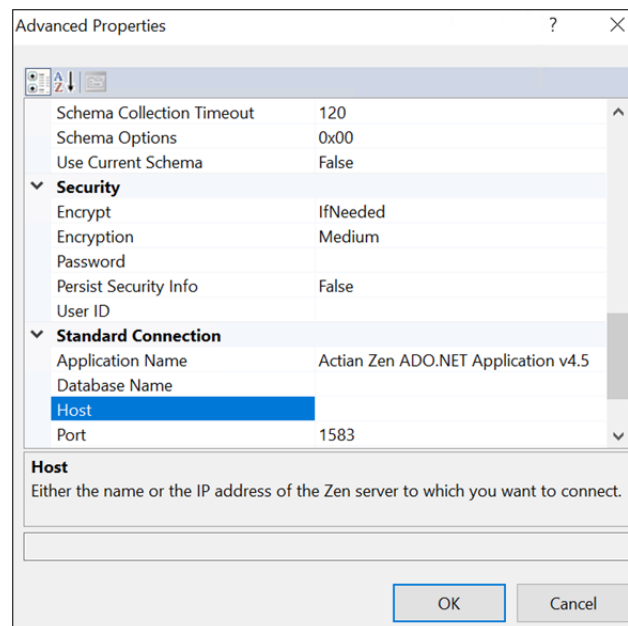
2. If the Zen data provider is displayed in the **Data source** field, skip to Step 4. Otherwise, click **Change**.
3. The Change Data Source window appears.



The 'Change Data Source' dialog box is shown. It has a title bar with a question mark and a close button. The main text says: 'Data source:' followed by a list box. The list box contains the following items: 'Actian Zen Database' (which is selected and highlighted in blue), 'Microsoft Access Database File', 'Microsoft ODBC Data Source', 'Microsoft SQL Server', 'Microsoft SQL Server Database File', 'Oracle Database', and '<other>'. To the right of the list box is a 'Description' section with the text: 'Select Actian Zen 100% managed provider for fast, reliable and secure access to Zen Databases.' Below the list box is a 'Data provider:' label and a dropdown menu showing 'Actian Zen ADO.NET Data Provider'. Below the dropdown menu is a checkbox labeled 'Always use this selection' which is unchecked. At the bottom of the dialog are two buttons: 'OK' and 'Cancel'.

- a. In the Data source list box, select **Actian Zen Database**.
- b. In the Data provider list, select **Actian Zen ADO.NET Data Provider**.
- c. If you want to use these selections for other connections, select the **Always use this selection** check box.

-
- d. Click **OK** to return to the Add Connection window.
 4. In the Add Connection window, do the following steps:
 - a. Enter the Host name.
 - b. Enter the User ID and password. These values are required for authentication.
 - c. (Optional) If you want to save the password for the lifetime of connection instance defined in Server Explorer, select the **Save my password** check box.
 - d. (Optional) In the Database entry field, enter the name of the database to which you want to connect.
 5. Click the **Advanced** button to specify additional provider-specific property values.



To change a value in the Advanced Properties dialog box, select or type the new value into the field and press ENTER. The value is added to the connection string that appears in the field below the description of the property. If you accept the default values, the connection string field remains unchanged. When you have made the necessary changes, click **OK** to return to the Add Connection window.

Advanced

EnableIPv6: Provides backward compatibility for connecting to a Zen server using an IPv4 address.

If set to True, a client with IPv6 protocol installed can connect to the server using either an IPv4 address or an IPv6 address.

If set to False, the clients run in the backward compatibility mode. The client always connects to the server using an IPv4 address.

The default value for 4.0 is set to True.

For more information about IPv6 formats, see [IPv6](#) in *Getting Started with Zen*.

Encoding: Type the ANSI name or Windows code page to be used for translating string data stored in the database. By default, the Windows code page is used.

Initial Command Timeout: Specifies the default wait time (timeout in seconds) before the data provider terminates the attempt to execute the command and generates an error. This option provides the same functionality as the PsqlCommand object's CommandTimeout property without the need to make changes to the application code. Subsequently, an application can use the CommandTimeout property to override the Initial Command Timeout connection string option.

The initial default value is 30 seconds.

Note: Set the Initial Command Timeout option to a value that is greater than the largest default deadlock detection and timeout value on the server. This ensures that the application receives a more meaningful reply in case of a timeout.

Initialization String: Type one statement that will be issued immediately after connecting to the database to manage session settings.

Example: To handle CHAR columns that are padded with NULLs, set the value to:

```
Initialization String=SET ANSI_PADDING ON
```

Note: If the statement fails to execute for any reason, the connection to the server fails. The data provider throws an exception that contains the errors returned from the server.

Parameter Mode: Select the behavior of native parameter markers and binding. This allows applications to reuse provider-specific SQL code and simplifies migration to Zen data providers. Note that this option does not apply to Zen ADO.NET Entity Framework data providers.

If set to ANSI (the default), the ? character is processed as a parameter marker and bound as ordinal.

If set to BindByOrdinal, native parameter markers are used and are bound as ordinal.

If set to BindByName, native parameter markers are used and are bound by name.

PVTranslate: Select whether the client should negotiate a compatible encoding with the server.

If set to Auto, then the data provider will set the Encoding connection property to the database code page. In addition, SQL query text will be sent to the engine using UTF-8 encoding instead of the data encoding. This preserves NCHAR string literals in the query text.

If set to Nothing (the default), the setting for the Encoding connection property is used.

Timestamp: Select whether Zen time stamps are stored and retrieved as strings.

If set to DateTime (the initial default), the data provider maps time stamps to DateTime. This setting may be appropriate when native precision is required, for example, when using the CommandBuilder with a time stamp.

If set to String, the data provider maps Zen time stamps as strings.

TimeType: Select whether Zen times are retrieved as Timespan or DateTime in Zen ADO.NET data providers.

If set to As DateTime, the data provider maps the SQL type TIME to the .NET type System.DateTime.

If set to As TimeSpan, the data provider maps the SQL type TIME to the .NET type System.DateTimespan.

Connection Pooling

Connection Reset: Select whether a connection that is removed from the connection pool for reuse by an application will have its state reset to the initial configuration settings of the connection.

If set to False (the initial default), the data provider does not reset the state of the connection.

Connection Pool Behavior: Select the order in which a connection is removed from the connection pool for reuse, based on how frequently or how recently the connection has been used.

If set to MostRecentlyUsed, the data provider uses a Last In First Out (LIFO) approach to return the connection that was returned to the pool most recently.

If set to LeastRecentlyUsed, the data provider uses a First In First Out (FIFO) approach to return the connection with the lowest use count. This value ensures a balanced use of connections in the pool.

If set to MostFrequentlyUsed, the data provider returns the connection with the highest use count. This value enables applications to give preference to the most seasoned connection.

If set to `LeastFrequentlyUsed`, the data provider returns the connection with the lowest use count. This value ensures a balanced use of connections in the pool.

Connection Timeout: Type the number of seconds after which the attempted connection to the server will fail if not yet connected. If connection failover is enabled, this option applies to each connection attempt.

If set to 0, the data provider never times out on a connection attempt.

The initial default is 15 seconds.

Load Balance Timeout: Type the number of seconds to keep connections in a connection pool. The pool manager periodically checks all pools, and closes and removes any connection that exceeds this value. The `Min Pool Size` option can cause some connections to ignore the value specified for the `Load Balance Timeout` option.

The value can be any integer from 0 to 65535.

If set to 0 (the initial default), the connections have the maximum timeout.

See [Removing Connections from a Pool](#) for a discussion of connection lifetimes.

Max Pool Size: Type the maximum number of connections within a single pool. When the maximum number is reached, no additional connections can be added to the connection pool.

The value can be any integer from 1 to 65535.

The initial default is 100.

Max Pool Size Behavior: Select whether the data provider can exceed the number of connections specified by the `Max Pool Size` option when all connections in the connection pool are in use.

If set to `SoftCap`, when all connections are in use and another connection is requested, a new connection is created, even when the connection pool exceeds the number set by the `MaxPoolSize` option. If a connection is returned and the pool is full of idle connections, the pooling mechanism selects a connection to be discarded so the connection pool never exceeds the `Max Pool Size`.

If set to `HardCap`, when the maximum number of connections allowed in the pool are in use, any new connection requests wait for an available connection until the `Connection Timeout` is reached.

Min Pool Size: Type the minimum number of connections that are opened and placed in a connection pool when it is created. The connection pool retains this number of connections, even when some connections exceed their `Load Balance Timeout` value.

The value can be any integer from 0 to 65535.

If set to 0 (the initial default), no additional connections are placed in the connection pool when it is created.

Pooling: Select True (the initial default) to enable connection pooling.

Failover

Alternate Servers: Type a list of alternate database servers to which the data provider will try to connect if the primary database server is unavailable. Specifying a value for this property enables connection failover for the data provider.

For example, the following Alternate Servers value defines two alternate servers for connection failover:

```
Alternate Servers="Host=AcctServer;Port=1584,  
Host=123.456.78.90;Port=1584"
```

Connection Retry Count: Type the number of times the data provider tries to connect to the primary server, and, if specified, the alternate servers after the initial unsuccessful attempt.

The value can be any integer from 0 to 65535.

If set to 0 (the initial default), there is no limit to the number of attempts to reconnect.

Connection Retry Delay: Type the number of seconds the data provider waits after the initial unsuccessful connection attempt before retrying a connection to the primary server, and, if specified, the alternate servers.

The initial default is 3.

This property has no effect unless the Connection Retry Count property is set to an integer value greater than 0.

Load Balancing: Select True or False to determine whether the data provider uses client load balancing in its attempts to connect to primary and alternate database servers.

If set to False (the initial default), the data provider does not use client load balancing.

Performance

Enlist: Select True or False to determine whether the data provider automatically attempts to enlist the connection in creating the thread's current transaction context.

Note: Because Zen does not support distributed transactions, any attempt to enlist the connection in the thread's current transaction context will fail.

If set to False (the initial default), the data provider does not automatically attempt to enlist the connection.

If set to True, the data provider returns an error on the connection if a current transaction context exists. If a current transaction context does not exist, the data provider raises a warning.

Max Statement Cache Size: Type the maximum number of statements generated by the application that can be held in the statement cache for this connection.

The value can be 0, or any integer greater than 1.

If set to 0, statement caching is disabled.

If set to an integer greater than 1, the value determines the number of statements that can be held in the statement cache.

The initial default is 10.

Statement Cache Mode: Select the statement caching mode for the lifetime of the connection. See [Using Statement Caching](#) for more information.

If set to Auto, statement caching is enabled. Statements marked as Implicit by the Command property StatementCacheBehavior are cached. These commands have a lower priority than that of explicitly marked commands, that is, if the statement pool reaches its maximum number of statements, the statements marked implicit are removed from the statement pool first to make room for statements marked Cache.

If set to ExplicitOnly (the initial default), only commands that are marked Cache by the StatementCacheBehavior property are cached. Note that this is the only valid value for Zen ADO.NET Entity Framework data providers.

Schema Information

Schema Collection Timeout: Type the number of seconds after which an attempted schema collection operation fails if it is not yet completed.

The initial default is 120.

Schema Options: Specifies additional database metadata that can be returned. By default, the data provider prevents the return of some performance-expensive database metadata to optimize performance. If your application needs this database metadata, specify the name or hexadecimal value of the metadata.

This option can affect performance.

If set to ShowColumnDefaults or 0x04, column defaults are returned.

If set to ShowParameterDefaults or 0x08, column defaults are returned.

If set to FixProcedureParamDirection or 0x10, procedure definitions are returned.

If set to ShowProcedureDefinitions or 0x20, procedure definitions are returned.

If set to ShowViewDefinitions or 0x40, view definitions are returned.

If set to ShowAll or 0xFFFFFFFF, all database metadata is returned.

For example, to return descriptions of procedure definitions, specify Schema Options=ShowProcedureDefinitions or Schema Options=0x20.

To show more than one piece of the omitted database metadata, specify either a comma-separated list of the names, or the sum of the hexadecimal values of the column collections that you want to restrict.

See [Date and Time Canonical Functions](#) for the name and hexadecimal value of the database metadata that the data provider can add.

Use Current Schema: This connection string option is not supported. Setting it will cause the data provider to throw an exception.

Security

Encrypt: Select whether the data provider uses Encrypted Network Communications, also known as wire encryption.

If set to IfNeeded (the initial default), the data provider reflects the server's setting.

If set to Always, the data provider uses encryption, or, if the server does not allow wire encryption, returns an error.

If set to Never, the data provider does not use encryption and returns an error if wire encryption is required by the server.

Encryption: Select the minimum level of encryption allowed by the data provider. The meaning of these values depends on the encryption module used. With the default encryption module, the values Low, Medium, and High correspond to 40-, 56-, and 128-bit encryption, respectively.

The initial default is Medium.

Password: Type a case-insensitive password used to connect to your Zen database. A password is required only if security is enabled on your database. If so, contact your system administrator to get your password.

Persist Security Info: Select whether to display secure information in clear text in the `ConnectionString` property.

If set to `True`, the value of the password connection string option is displayed in clear text.

If set to `False` (the initial default), the data provider does not display secure information in clear text.

User ID: Type the default Zen user name used to connect to your Zen database.

Standard Connection

Database Name: Type a string that identifies the internal name of the database to which you want to connect.

If you enter a value for this field, the `Server DSN` field is not available.

Host: Type the name or the IP address of the Zen server to which you want to connect. For example, you can specify a server name such as `accountingserver`. Or, you can specify an IPv4 address such as `199.262.22.34` or an IPv6 address such as `2001:DB8:0000:0000:8:800:200C:417A`.

Port: Type the TCP port number of the listener running on the Zen database.

The default port number is 1583.

Server DSN: The name of the data source on the server, such as `DEMODATA`.

If you enter a value for this field, the `Database Name` field is not available.

Tracing

Enable Trace: Type a value of 1 or higher to enable tracing. If set to 0 (the default), tracing is not enabled.

Trace File: Type the path and name of the trace file. If the specified trace file does not exist, the data provider creates it. The default is an empty string.

1. Click **Test Connection**. At any point during the configuration process, you can click **Test Connection** to attempt to connect to the data source using the connection properties specified in the Add Connection window.

-
- If the data provider can connect, it releases the connection and displays a **Connection Established** message. Click **OK**.
 - If the data provider cannot connect because of an incorrect environment or incorrect connection value, it displays an appropriate error message.

2. Click **OK**.

Note: If you are configuring alternate servers for use with the connection failover feature, be aware that the Test Connection button tests only the primary server, not the alternate servers.


3. Click **OK** or **Cancel**. If you click **OK**, the values you have specified become the defaults when you connect to the data source. You can change these defaults by using this procedure to reconfigure your data source. You can override these defaults by connecting to the data source using a connection string with alternate values.

Adding Connections with the Data Source Configuration Wizard

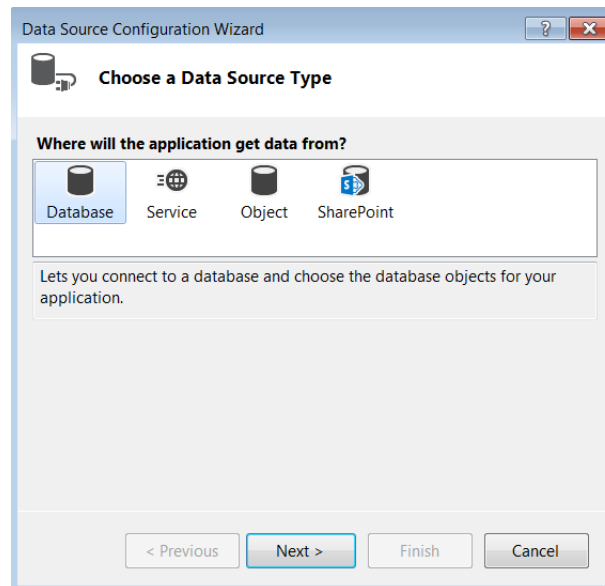
You can add a new connection to your application using the Data Configuration Wizard.

To add a connection

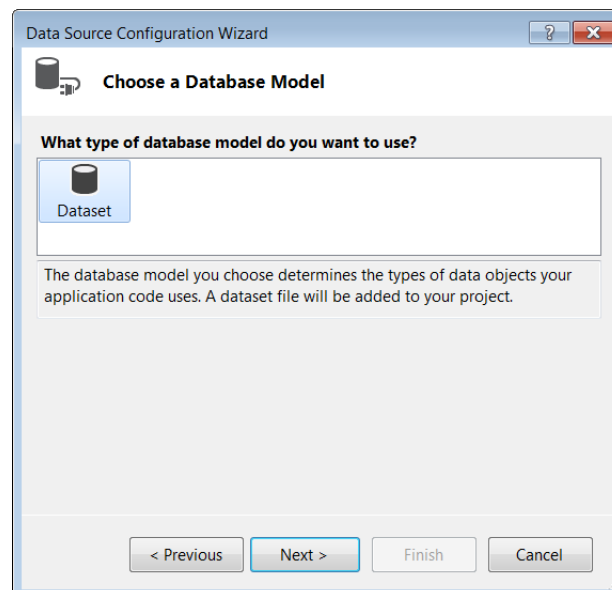
1. In the Data Sources window in Visual Studio, select **Add New Data Source**. To open the Data Sources window, select **View** from the main menu and then select **Other Windows > Data Sources**.




The Data Source Configuration Wizard appears.



2. Select **Database** and then click **Next**. The Choose Your Data Connection window appears.



-
3. Click **New Connection** The Add Connection window is displayed. Continue from Step in [Adding Connections in Server Explorer](#).



Using the Zen Performance Tuning Wizard

The Zen Performance Tuning Wizard leads you step-by-step through a series of questions about your application. Based on your answers, the wizard provides the optimal settings for performance-related connection string options for your Zen data provider.

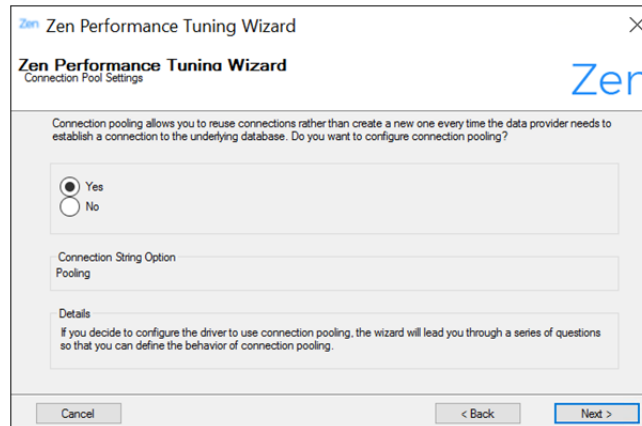
When you launch the Zen Performance Tuning Wizard from Visual Studio, you can:

- Generate the values for connection string options that are related to performance. These values can be copied into a connection string.
- Modify an existing connection.
- Generate a new application preconfigured with a connection string optimized for your environment. The Performance Tuning Wizard provides options to select the type of application and the version of ADO.NET code that you want to use.

To use the Zen Performance Tuning Wizard in Visual Studio

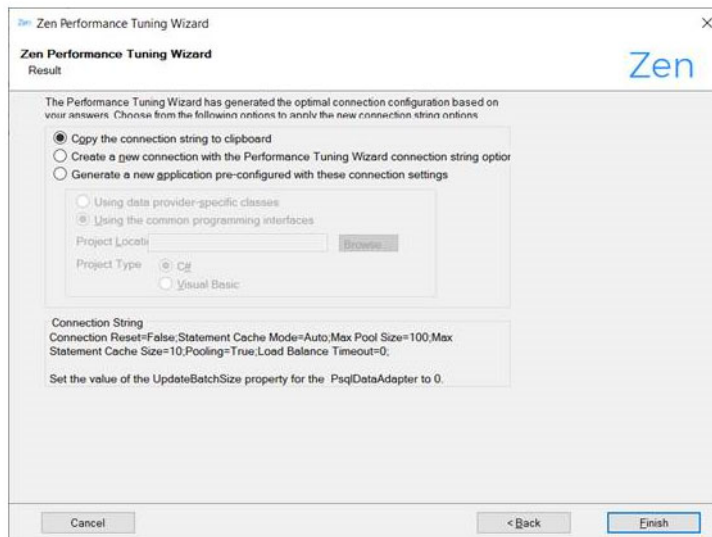
1. Do one of the following steps to start the Performance Tuning Wizard:
 - To create a new connection, select **Tools > Actian Zen > Run Zen Performance Tuning Wizard**. When the Performance Tuning Wizard Welcome dialog appears, click **Next**. Continue at Step 2.
 - To modify an existing connection, in Server Explorer, right-click a data connection, and then select **Run Zen Performance Tuning Wizard**. When the Performance Tuning Wizard Welcome dialog appears, click **Next**. Continue at Step 2.
2. The wizard presents a series of questions about your environment. Accept the default or change the answers as required and then, click **Next** to proceed, until you reach the Result page.

The following screen shot shows an example of one of the questions you may be asked.



3. When you have answered all questions for a data provider, the Result dialog appears, and a connection string is displayed.

The following screen shot shows the connection string options related to performance that the Zen Performance Tuning Wizard generated.



4. Select one of the following:
 - To make the connection string available to other applications (the initial default), select **Copy the connection string to clipboard**. You can use the connection string in other applications.
 - Based on whether you have used a new connection or an existing connection to launch the wizard, select one of the following options:

-
- **Create a new connection with the Performance Tuning Wizard connection string options**

When you select this option and click **Finish**, the Modify Connection dialog box appears, where you must specify the connection information, such as a host, password, user ID, and other information.

- **Reconfigure the connection with the additional Performance Tuning Wizard connection string options.**
- To create a new application, select **Generate a new application preconfigured with these connection settings.**

When you select this option and click **Finish**, a Zen application is generated using the Zen Project template. See [Creating a New Project](#) for more information about the provider-specific templates.

5. Define additional information for the new application:

- Select **Using data provider-specific interfaces** to create an application compatible with the ADO.NET 2.0 specification.
- Select **Using common programming interfaces** to create an application that uses the ADO.NET common programming model.
- Type the location for the project, or click **Browse** to select the location.
- Select the project type. By default, the Wizard creates a C# application.

6. Click **Finish** to exit the Zen Performance Tuning Wizard.

Using Provider-Specific Templates

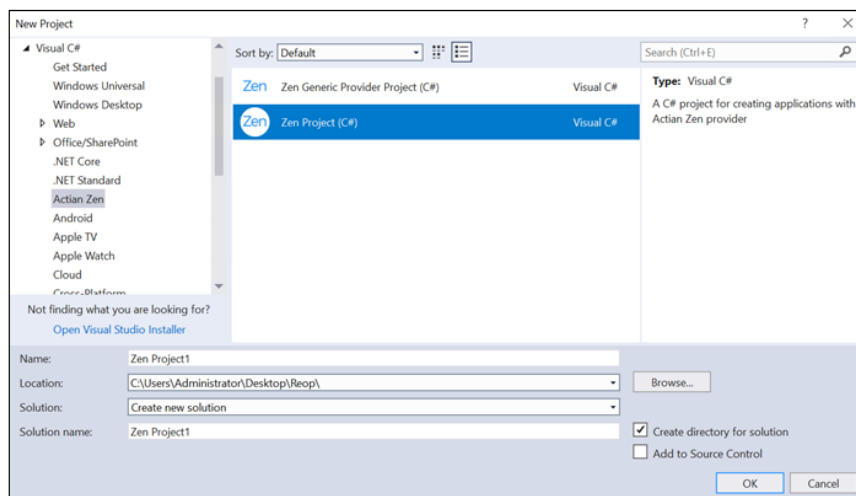
Visual Studio offers a set of templates to help you build applications that automatically include features such as SQL leveling.

Creating a New Project

When you create a new project in Visual Studio, you can use a template specific to the Zen data provider, or a template that creates an application with generic code.

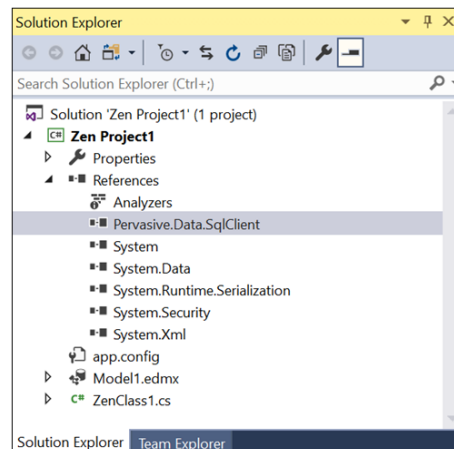
In the following example, we create a new project in Visual Studio using the template for the Zen data provider.

1. Select **File > New > Project** The New Project dialog appears.
2. In the Installed List, select **Visual C# > Actian Zen**.
3. Select **Zen Project** in the middle pane.



4. Make changes to the other fields if required, and click **OK**.

5. The new project appears in the Solution Explorer. The namespace for the Zen ADO.NET data provider is automatically added to the project.

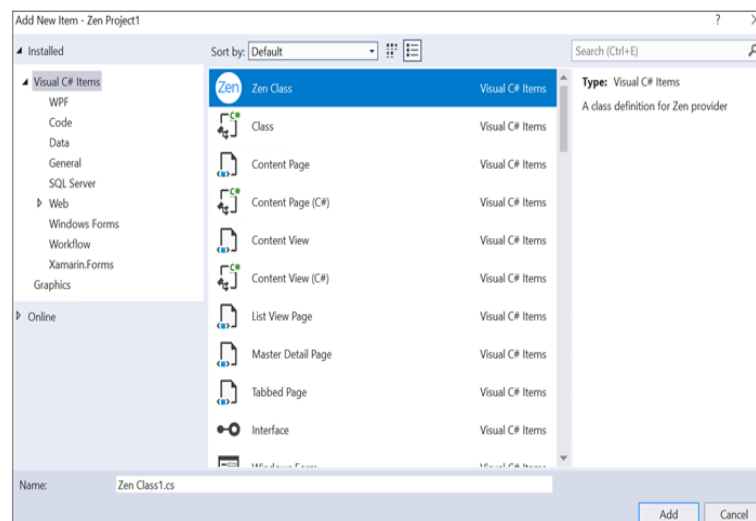


Note: If you are using the ADO.NET 2.0 common programming model, select the Zen Generic Provider Project template. In this case, the project does not require a specific reference to an assembly.

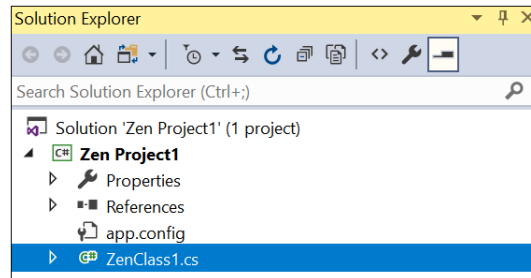
Adding a Template to an Existing Project

To add a Zen template to an existing project

1. In Solution Explorer, right-click the project and select **Add > New Item**.
2. In the **Add New Item** dialog, select the Zen class.



-
3. Click **Add**. The class for the Zen data provider is added to the project.



Using the Zen Visual Studio Wizards

Wizards simplify typical tasks that you perform when you create an application:

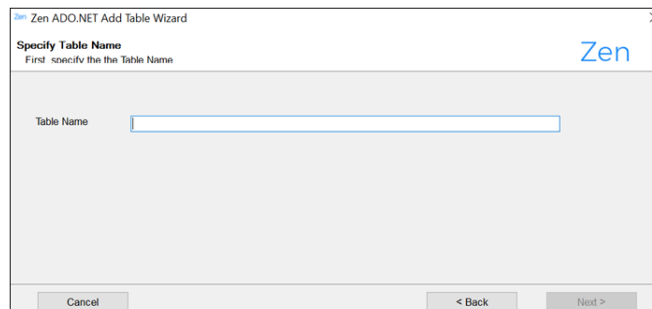
- [Creating Tables with the Add Table Wizard](#)
- [Creating Views with the Add View Wizard](#)

Before beginning this procedure, create a project using a Zen template, as described in [Creating a New Project](#), and add a data connection.


Creating Tables with the Add Table Wizard

You can quickly and easily define new tables in Visual Studio using the Zen ADO.NET Add Table Wizard.

1. Select **View > Server Explorer**.
2. Double-click a data source connection to expose the nodes below it.
3. Right-click the **Tables** node, and select **Add New Table**. The Zen ADO.NET Add Table Wizard appears.
4. Click **Next**. The Specify Table Name dialog appears.
5. In the **Table Name** field, type a name for the table.



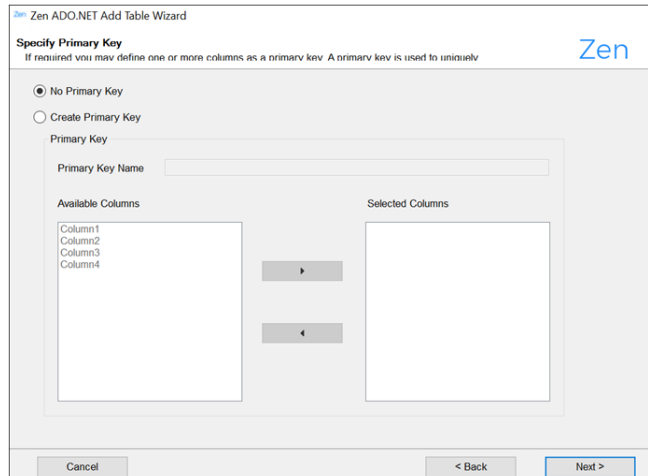
6. Click **Next**. The Specify Column(s) dialog appears.



7. Define the columns for the new table. Your choices may cause additional fields to appear in the Data Type Options pane.

- Click **Add** to add a column to the table. The Column Name and Data Type fields become editable.
- Type a name in the Column Name field.
- Select the data type for the column, and, if required, supply any additional information:
 - If you select a character data type, the Length field appears in the Data Type Options pane. Type the maximum size of the column (in bytes).
 - If you select Number, the Precision and Scale fields appear in the Data Type Options pane.
- If the column can have a Null value, select the **Allow Null** check box.
- To remove a column from the table, select the column name and then click **Remove**.

8. Click **Next**. The Specify Primary Key dialog appears.

The image shows the 'Specify Primary Key' dialog box from the 'Zen ADO.NET Add Table Wizard'. The dialog has a title bar with the Zen logo and a close button. Below the title bar, it says 'Specify Primary Key' and 'If required you may define one or more columns as a primary key. A primary key is used to uniquely'. There are two radio buttons: 'No Primary Key' (selected) and 'Create Primary Key'. Below the 'Create Primary Key' option, there is a 'Primary Key' section with a 'Primary Key Name' text box. Below that, there are two list boxes: 'Available Columns' (containing Column1, Column2, Column3, Column4) and 'Selected Columns' (empty). Between the list boxes are two arrow buttons (right and left). At the bottom, there are three buttons: 'Cancel', '< Back', and 'Next >'.

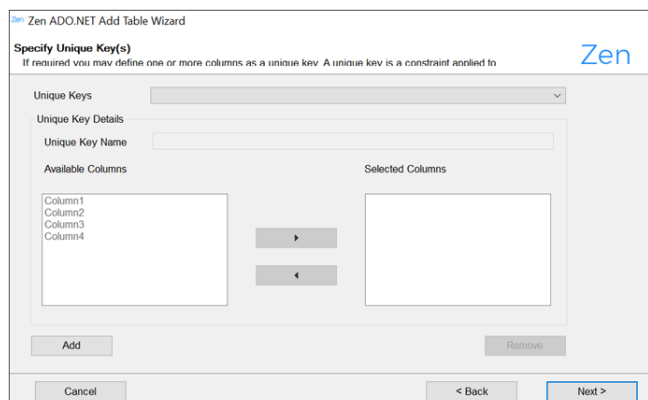
9. Do one of the following:

- If you do not want to specify a primary key for the table, select **No Primary Key**, and then click **Next**. The Specify Unique Key(s) dialog appears. Continue at Step 12.
- If you want to specify a primary key for the table, select **Create Primary Key**, and then continue at Step 10.

10. Complete the fields on the Specify Primary Key dialog:

- In the Primary Key Name field, type the name for the primary key, or accept the default name.
- Select a column from the Available Columns field and move it to the Selected Columns field.

11. Click **Next**. The Specify Unique Key(s) dialog appears.

The image shows the 'Specify Unique Key(s)' dialog box from the 'Zen ADO.NET Add Table Wizard'. The dialog has a title bar with the Zen logo and a close button. Below the title bar, it says 'Specify Unique Key(s)' and 'If required you may define one or more columns as a unique key. A unique key is a constraint applied to'. There is a 'Unique Keys' dropdown menu. Below it, there is a 'Unique Key Details' section with a 'Unique Key Name' text box. Below that, there are two list boxes: 'Available Columns' (containing Column1, Column2, Column3, Column4) and 'Selected Columns' (empty). Between the list boxes are two arrow buttons (right and left). Below the list boxes are two buttons: 'Add' and 'Remove'. At the bottom, there are three buttons: 'Cancel', '< Back', and 'Next >'.

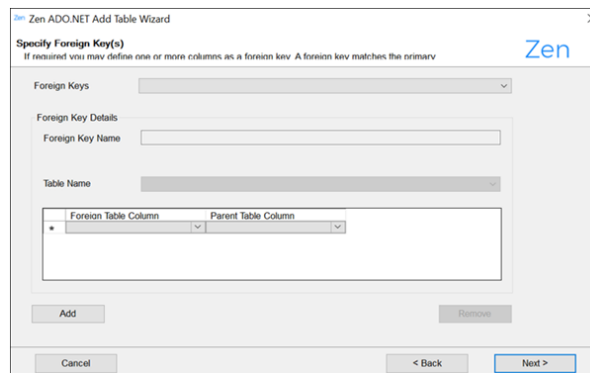
12. Do one of the following:

- If you do not want to specify unique keys for the table, click **Next**. The Specify Foreign Key(s) dialog appears. Continue at Step 15.
- If you want to specify one or more unique keys for the table, continue at Step 13.

13. Click **Add**. The fields on the dialog become selectable:

- In the **Unique Keys** drop-down list, select a unique key.
- In the **Unique Key Name** field, edit the name or accept the default name.
- In the **Available Columns** list box, select one or more columns to be used to specify the unique key, and move them to the Selected Columns list box.

14. Click **Next**. The Specify Foreign Key(s) dialog appears.




15. Do one of the following:

- If you do not want to specify foreign keys for the table, click **Next**. The Review SQL dialog appears. Continue at Step 18.
- If you want to specify one or more foreign keys for the table, continue at Step 16.

16. Click **Add**. The fields on the dialog become selectable:

- In the **Foreign Keys** drop-down list, select a foreign key.
- In the **Foreign Key Name** field, edit the name or accept the default name.
- In the **Table Schema** list, select a table schema.
- In the **Table Name** list, select a table schema.
- In the **Foreign Key Column** list, select one or more columns to be used to specify the foreign key.
- In the **Parent Table Column** list, select the corresponding column from the parent table.

17. Click **Next**. The Review SQL dialog appears.



18. Review the SQL statement that has been generated by your choices.

- If you are satisfied with the SQL statement, click **Finish**. The table that you created appears in Server Explorer under the Tables node for the connection.
- If you want to supplement the SQL statement, for example, add a view or specific keywords, continue at Step 19.

19. Select the **Edit SQL** check box. The text in the Generated SQL field becomes editable.

Note: When you select the Edit SQL check box, the Back button is disabled.


20. When you are satisfied with your changes to the SQL statement, click **Finish**. The table that you created appears in Server Explorer under the Tables node for the connection.

Creating Views with the Add View Wizard


You can quickly and easily define new views in Visual Studio using the Zen Add View Wizard.

1. Select **View > Server Explorer** if it is not already open.
2. Double-click a data source connection to expose the nodes under it.
3. Right-click the **Views** node, and select **Add New View**. The Zen Add View Wizard welcome dialog appears.
4. Click **Next**. The Specify View Name dialog appears.


5. Type a name for the view in the View Name field.



6. Click **Next**. The Select Table(s) and/or Column(s) dialog appears.



7. In the **List of Tables and columns** list box, select the tables or columns that will make up the view, and move them to the Selected Columns column.
8. Click **Next**. The Review SQL dialog appears.



9. Review the SQL statement that has been generated by your choices.

-
- If you are satisfied with the SQL statement, click **Finish**. The view that you created appears in Server Explorer under the Views node for the connection.
 - If you want to supplement the SQL statement, for example, add a view or specific keywords, continue at Step 11.
10. Select the **Edit SQL** check box. The text in the Generated SQL field becomes editable.
Note: When you select the Edit SQL check box, the Back button is disabled.
 11. When you are satisfied with your changes to the SQL statement, click **Finish**. The view that you created appears in Server Explorer under the Views node for the connection.

Adding Components from the Toolbox

You can add components from the Visual Studio Toolbox to a Windows Forms application. For information about creating Windows Forms applications, refer to the Visual Studio online Help.

Before beginning this procedure, create a Windows Forms application and add a data connection.

To add Zen data provider components to a Windows Forms application

1. Select **View > Toolbox**. Scroll down the Toolbox until the Zen ADO.NET Provider section appears.
2. Select the **PsqlCommand** widget and drag it onto the Windows Forms application.
3. Continue adding widgets to the application as needed.

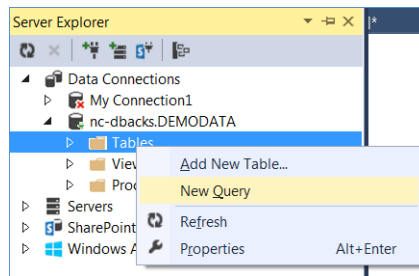
Data Provider Integration Scenario

Because the Zen data provider is integrated into Visual Studio, many typical data access tasks can be simplified. For example, after making the connection to the database, you can create queries using Query Builder.


The Query Builder can help you graphically design database queries.

To create a simple query

1. Establish a data source connection (see [Using the Zen Visual Studio Wizards](#)).
2. Select the data source in Server Explorer.
3. Right-click **Tables** and select **New Query**.

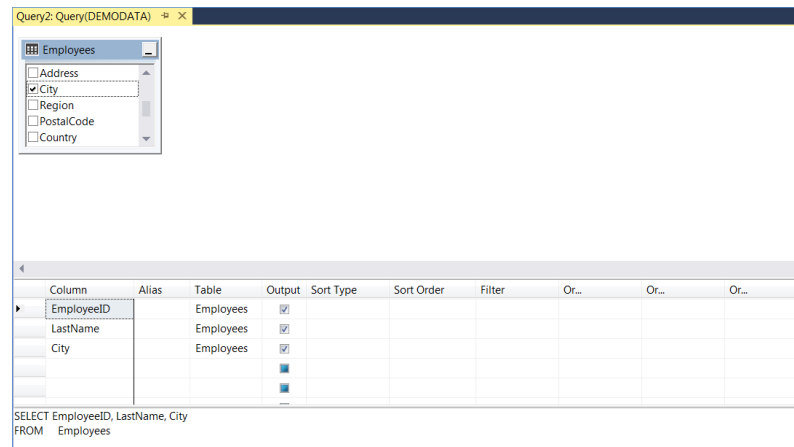


4. The Add Table window appears. Select the table that contains the data that you want to use; then, click **Add**.




5. Click **Close** to close the Add Table window.

-
6. Select the columns that you want returned. In this example, we select the id, name, and salary columns of the employee table.



7. Click the **Execute SQL** button on the toolbar.



8. Examine the results displayed.

Using the Microsoft Enterprise Libraries

Using the Microsoft Enterprise Libraries can simplify application development by wrapping common tasks, such as data access, into portable code that makes it easier to move your application from one DBMS to another.

The Zen data providers support the Data Access Application Blocks (DAAB) and Logging Application Blocks (LAB). The classes in the DAABs provide access to the most frequently used features of ADO.NET. Applications can use the DAABs for tasks such as passing data through application layers and returning changed data back to the database. Using DAABs eliminates the need to keep writing the same data access tasks for each new or revised application, so you can spend your time more productively

Applications that use the standard Logging Application Block and design patterns can quickly display the SQL that is generated as part of Zen ADO.NET Entity Framework data providers.

To use features of the Enterprise Library with your data provider, download Microsoft Enterprise Library 6.0 from [https://docs.microsoft.com/en-us/previous-versions/msp-n-p/dn169621\(v=pandp.10\)](https://docs.microsoft.com/en-us/previous-versions/msp-n-p/dn169621(v=pandp.10)). You can also download the Enterprise Library documentation, which contains detailed information about using the application blocks.

Data Access Application Blocks

The Data Access Application Block (DAAB) is designed to allow developers to replace ADO.NET boiler-plate code with standardized code for everyday database tasks. The overloaded methods in the Database class can:

- Return Scalar values or XmlReaders
- Determine which parameters it needs and create them
- Involve commands in a transaction

If your application needs to address specific DBMS functionality, you can use the ADO.NET Data Provider.

When Should You Use the DAAB?

The DAAB includes a small number of methods that simplify the most common methods of accessing a database. Each method encapsulates the logic required to retrieve the data and manage the connection to the database. You should consider using the application block if your application uses standard data access techniques.

The DAAB is used with ADO.NET, increasing efficiency and productivity when creating applications for ADO.NET. The abstract Database class provides a number of methods, such as `ExecuteNonQuery`, `ExecuteReader`, and `ExecuteScalar`, that are the same as the methods used by the `DbCommand` class, or, if you are using database-specific code, a data provider-specific class such as `PsqlCommand`.

Although using the default DAAB during development is convenient, the resulting application lacks portability. When you use the provider-specific Zen DAAB implementation, the application includes the ADO.NET Data Provider's SQL leveling capabilities. You have more flexibility, whether your application needs to access multiple databases, or whether you anticipate a change in your target data source.

Should You Use Generic or Database-specific Classes?

The application block supplements the code in ADO.NET that allows you to use the same code with different database types. You have two choices when using the DAAB with the ADO.NET data provider:

- The `GenericDatabase` class
- The provider-specific Zen DAAB implementation

The GenericDatabase class option is less suited to applications that need specific control of database behaviors. For portability, the provider-specific solution is the optimal approach.

If your application needs to retrieve data in specialized way, or if your code needs customization to take advantage of features specific to Zen, using the ADO.NET Data Provider might suit your needs better.

Configuring the DAAB

Before you can configure the DAAB for use with your application, you must set up the environment:

1. Download and install Microsoft Enterprise Library 6.0 from [https://docs.microsoft.com/en-us/previous-versions/msp-n-p/dn169621\(v=pandp.10\)](https://docs.microsoft.com/en-us/previous-versions/msp-n-p/dn169621(v=pandp.10)).
2. Open the Zen DAAB project for your data provider, located in *install_dir\Enterprise Library\Src\CS\Actian Zen*.
3. Then, compile the project and note the output directory.

Configuring the Data Access Application Block consists of two procedures:

- [Adding a New DAAB Entry](#)
- [Adding the Data Access Application Block to Your Application](#)

Adding a New DAAB Entry

Now, use the Enterprise Library Configuration Tool to add a new DAAB entry:

1. Right-click **Enterprise Library Configuration**, and select **New Application**.
2. Right-click **Application Configuration**, then select **New > Data Access Application Block**. The Enterprise Library Configuration window appears.
3. In the Name field, type a name for the DAAB connection string.
4. In the ConnectionString field, enter a connection string.
5. Right-click the field, and select the data provider. For example, select `Pervasive.Data.SqlClient`.
6. Right-click **Custom Provider Mappings** and select **New > Provider Mappings**.
7. In the Name field, select the data provider name you specified in Step 5.

-
8. Select the **TypeName** field, and then choose the **Browse** button to navigate to the Debug output directory of the Zen DAAB that you want to build. Then select the TypeName. For example, the Zen TypeName is Pervasive.EnterpriseLibrary.Data.Pervasive.dll.
 9. Leave the Enterprise Library Configuration window open for now and do not save this configuration until you complete the following section.

Adding the Data Access Application Block to Your Application

To add the DAAB to a new or existing application, perform these steps:

1. Add two additional References to your Visual Studio solution:

- Enterprise Library Shared Library
- Enterprise Library Data Access Application Block

2. Add the following directive to your C# source code:

```
using Microsoft.Practices.EnterpriseLibrary.Data;
```

```
using System.Data;
```

3. Rebuild the solution to ensure that the new dependencies are functional.
4. Determine the output Debug or Release path location of your current solution, and switch back to the Enterprise Library Configuration window (see [Adding a New DAAB Entry](#)).
5. Right-click the connection string under the **Application Configuration** node and select Save Application.
6. Navigate to the Debug or Release output directories of your current solution, and locate the .exe file of the current solution.
7. Click the file name once, and add .config to the name, for example MyApp.config.
8. Ensure that the Save as type All Files is selected, and click **Save**.
9. Using File Explorer, copy the Pervasive.EnterpriseLibrary.Data.Pervasive.dll from the Zen DAAB directory.
10. Place the copy of this DLL into either the Debug or Release output directory of your current solution.

Using the DAAB in Application Code

Now that you have configured the DAAB, you can build applications on top of this DAAB.

In the example below, we use the DAAB MyApp and the DatabaseFactory to generate an instance of a Database object backed by a Zen data source.

```
using System;
using System.Collections.Generic;
using System.Text;
using Microsoft.Practices.EnterpriseLibrary.Data;
using System.Data;

namespace DAAB_Test_App_1
{
    class Program
    {
        static void Main(string[] args)
        {
            Database database = DatabaseFactory.CreateDatabase("MyApp");
            DataSet ds = database.ExecuteDataSet(CommandType.TableDirect,
"SQLCOMMANDTEST_NC_2003SERVER_1");
        }
    }
}
```

The Microsoft Enterprise Library DAAB coding patterns are now at your disposal.

Logging Application Blocks

Using the Enterprise Library Logging Application Block (LAB) makes it easier to implement common logging functions. Zen ADO.NET Entity Framework data providers use the standard Logging Application Block and design patterns, and offer LAB customizations for additional functionality.

When Should You Use the LAB?

If your applications have a requirement to log information to a database, a message queue, the Windows Event Log, or Windows Management Instrumentation (WMI), the LAB provides this functionality. In particular, the LAB is useful if you need to filter logging messages based on category or priority, if you need to format the messages, or if you need to change the destination of the message without changing the application code.

Configuring the LAB

A logging capability can be added to an application by adding an entry to an applications configuration file (either app.config or web.config) using the Enterprise Library configuration tool. This tool contains specific instructions in order to enable the Logging Application Block .config file. The tool also contains the necessary AppSetting to enable the LAB.

To enable the Logging Application Block output, set the environment property `Psql_Enable_Logging_Application_Block_Trace` to true. Alternatively, in the app.config file, set the AppSetting property `Psql.EnableLoggingApplicationBlock` to true.

A following C# code snippet shows the loggingConfiguration property of the app.config file.

```
<loggingConfiguration name="Logging Application Block" tracingEnabled="true"
defaultCategory="General" logWarningsWhenNoCategoriesMatch="true">
```

Setting either of these properties to false disables the logging block.

If enabled, the data provider must establish a new `LogEntry` entry instance for each SQL statement generated by the ADO.NET Entity Framework canonical query tree.

The SQL logged to the Logging Block must be the SQL that is ultimately transmitted to over the wire.

To configure the Logging Application Block

1. Select **Start > Programs > Microsoft patterns and practices > Enterprise Library 6.0 > Enterprise Library Configuration**. The Enterprise Library Configuration window appears.

-
2. Select **File > New Application**.
 3. Right-click the **Application Configuration** node and select **New > Logging Application Block**.
 4. Right-click **Category Sources**, and select **New > Category**.
 5. In the Name pane, select **Name**. Type the name of the new category, and then press **Enter**.
 6. From the SourceLevels list, set the logging level for the new category. By default, all logging levels are enabled.
 7. Right-click the new category and select **New > TraceListener Reference**. A Formatted EventLog TraceListener node is added. From the ReferencedTraceListener list, select **Formatted EventLog TraceListener**.
 8. Repeat steps 4 through 7 to create the following categories:
 - Zen Information: Information not related to errors
 - Zen Command: Enables SQL, Parameter, and DbCommandTree logging
 9. Select **File > Save Application**, and enter a name for your configuration file. By default, the file is saved to C:\Program Files\Microsoft Enterprise Library\Bin\filename.exe.config, where *filename* is the name you entered.

Adding a New Logging Application Block Entry

Now, use the Enterprise Library Configuration Tool to add a new Logging Application Block entry:

1. Start **Enterprise Library Configuration**, and select **File > New Application**.
2. Right-click **Application Configuration**, then select **New > Logging Application Block**. The Configuration section appears in the right pane.
3. In the **TracingEnabled** field, enter **True**.
4. Save the Logging application block.

Using the LAB in Application Code

The LAB that you configured must be added to the app.config or web.config file for your application.

The following settings can be used to enable and configure data provider interaction with the LAB.

- **EnableLoggingApplicationBlock:** Enables the Logging Application Block.
- **LABAssemblyName:** Specifies the assembly name to which the Logging Application Block applies.

Note: If you use any version of the LAB other than the Microsoft Enterprise Library 4.1 binary release, you must set the **LABAssemblyName**. For example, if you use an older or newer version of the LAB, or a version that you have customized, you must specify a value for **LABAssemblyName**.

- **LABLoggerTypeName:** Specifies the type name for the Logging Application Block.
- **LABLogEntryTypeName:** Specifies the type name for the LogEntry object.

The following code fragment provides an example of a Logging Application Block that could be added to a Zen data access application.

```
<loggingConfiguration name="Logging Application Block"
  tracingEnabled="true"
  defaultCategory="" logWarningsWhenNoCategoriesMatch="true">
  <listeners>
    <add fileName="rolling.log"
      footer="-----"
      header="-----"
      rollFileExistsBehavior="Overwrite"
      rollInterval="None" rollSizeKB="0"
      timeStampPattern="yyyy-MM-dd"

listenerDataType="Microsoft.Practices.EnterpriseLibrary.Logging.Configuration.RollingFlatFileTraceLi
stenerData, Microsoft.Practices.EnterpriseLibrary.Logging, Version=6.0.0.0, Culture=neutral,
PublicKeyToken=31bf3856ad364e35"
      traceOutputOptions="None" filter="All"
type="Microsoft.Practices.EnterpriseLibrary.Logging.TraceListeners.RollingFlatFileTraceListener,
Microsoft.Practices.EnterpriseLibrary.Logging, Version=6.0.0.0, Culture=neutral,
PublicKeyToken=31bf3856ad364e35"
      name="Rolling Flat File Trace Listener" />
    </listeners>
  <formatters>
    <add template="Message: {message}&#xD;&#xA;Category: {category}&#xD;&#xA;Priority:
{priority}&#xD;&#xA;EventId: {eventid}&#xD;&#xA;Severity:
{severity}&#xD;&#xA;Title:{title}&#xD;&#xA;&#xD;&#xA;"
      type="Microsoft.Practices.EnterpriseLibrary.Logging.Formatters.TextFormatter,
Microsoft.Practices.EnterpriseLibrary.Logging, Version=6.0.0.0, Culture=neutral,
PublicKeyToken=31bf3856ad364e35"
      name="Text Formatter" />
    </formatters>
  <categorySources>
    <add switchValue="All" name="Psql">
      <listeners>
        <add name="Rolling Flat File Trace Listener" />
      </listeners>
    </add>
  </categorySources>
  <specialSources>
    <allEvents switchValue="All" name="All Events" />
    <notProcessed switchValue="All" name="Unprocessed Category" />
  </specialSources>
</loggingConfiguration>
```

```
<errors switchValue="All" name="Logging Errors & Warnings">
  <listeners>
    <add name="Rolling Flat File Trace Listener" />
  </listeners>
</errors>
</specialSources>
</loggingConfiguration>
```

Additional Resources

The following sources provide additional information about using Application Blocks:

- The Microsoft patterns & practices Developer Center includes a topic on [Application Blocks](#).
- See "[Introduction to the Data Access Application Block](#)" for an overview of tasks and applications using data access application blocks.
- See Microsoft patterns & practices Enterprise Library topics on using logging application blocks.

A. .NET Objects Supported

ADO.NET 2.0 introduced a new set of classes that provided an additional, more generic interface between applications and data sources.

Predecessors of ADO.NET 2.0 opted for a tighter factoring of data providers into each specific instance of the data provider used by an application. In contrast, ADO.NET 2.0 and higher deliver a set of base classes that permit applications to handle a heterogeneous set of data sources with a single API, much like is done with ODBC and JDBC today. This means that in ADO.NET 2.0 and higher, all data classes derive from base classes, and exist in a specific dedicated namespace, `System.Data.Common`.

The data provider supports:

- [.NET Base Classes](#)
- [Data Provider-Specific Classes](#)
- [Zen Common Assembly Classes](#)

.NET Base Classes

The interfaces on which ADO.NET 1.0 and ADO.NET 1.1 data providers were built were retained for application compatibility. The base classes of ADO.NET 2.0 and higher provide additional functionality:

- DbCommand
- DbCommandBuilder
- DbConnection
- DbDataAdapter
- DbDataReader
- DBDataPermission
- DbParameter
- DbParameterCollection
- DbConnectionStringBuilder
- DbTransaction

From a day-to-day programming perspective, these classes are provided as abstract implementation. This means they cannot be instantiated directly, but must be used with Provider factories. Each data provider must supply a Factory class, such as PsqlFactory, that derives from the DbFactory class, which contains a set of static methods. Each of these static methods is a factory method for producing an instance of the base classes.

When a data provider is installed, it is registered with the .NET Framework. This allows the common .NET Framework DbFactory to locate any registered data provider that an application requires and provide a common mechanism to establish a connection to a data source. Ultimately, the .NET Framework provides a fully fledged common programming API for ADO.NET data sources.

Data Provider-Specific Classes

The ADO.NET Data Provider supports all of the .NET public objects. The ADO.NET Data Provider attaches the provider-specific prefix *Psql* to the public .NET objects, for example, *PsqlCommand*.

The following objects are described:

- [PsqlBulkCopy](#)
- [PsqlBulkCopyColumnMapping](#)
- [PsqlBulkCopyColumnMappingCollection](#)
- [PsqlCommand Object](#)
- [PsqlCommandBuilder Object](#)
- [PsqlConnection Object](#)
- [PsqlConnectionStringBuilder Object](#)
- [PsqlCredential Object](#)
- [PsqlDataAdapter Object](#)
- [PsqlDataReader Object](#)
- [PsqlError Object](#)
- [PsqlErrorCollection Object](#)
- [PsqlException Object](#)
- [PsqlFactory Object](#)
- [PsqlInfoMessageEventArgs Object](#)
- [PsqlParameter Object](#)
- [PsqlParameterCollection Object](#)
- [PsqlTrace Object](#)
- [PsqlTransaction Object](#)

For more information on public objects, refer to the Microsoft .NET Framework Version 2.0 SDK documentation.

PsqlBulkCopy

The PsqlBulkCopy object uses an API pattern similar to the ADO.NET Bulk API patterns, and has no provider-specific properties or methods. For information about the properties and methods supported, see the data provider online help and the Microsoft .NET Framework SDK documentation.

PsqlBulkCopyColumnMapping

The PsqlBulkCopyColumnMapping object uses an API pattern similar to the ADO.NET Bulk API patterns, and has no provider-specific properties or methods. For information about the properties and methods supported, refer to the data provider's online help and the Microsoft .NET Framework SDK documentation.

PsqlBulkCopyColumnMappingCollection

The PsqlBulkCopyColumnMappingCollection object follows an API pattern similar to the Microsoft SqlBulkCopyColumnMappingCollection class, and has no provider-specific properties or methods. For information about the properties and methods supported, refer to the data provider's online help and the Microsoft .NET Framework SDK documentation.

PsqlCommand Object

The following table describes the public properties of the PsqlCommand object.

Property	Description
AddRowID	Adds the ROWID as part of the Select list of a SQL statement. If set to true, the values returned in the ROWID column are used to generate more efficient Insert, Delete, and Update commands when using the PsqlCommandBuilder. If set to false (the initial default), the data provider does not add the ROWID column to the Select list.
ArrayBindCount	Specifies the number of rows of parameters that will be used. The application must set this property before executing a command that uses parameter array binding. The count must equal the length of each of the arrays that is set for each parameter value. The initial default value is 0. The application does not use parameter array binding.

Property	Description
ArrayBindStatus	<p>Returns an array of row status values. This property enables the application to inspect the per row status after executing a command that uses parameter array binding. The property's type is an array of <code>PsqlRowStatus</code>.</p> <p>Parameter array binding is performed as a single atomic operation. This means that if the operation succeeds, every entry will be set to OK; if the operation fails, none of the entries will be set to OK.</p> <p>The <code>PsqlRowStatus</code> enumeration has the following possible values:</p> <ul style="list-style-type: none"> • OK. The operation succeeded. All entries are marked as OK. • Failed. The operation failed. The data provider assigns this value to all status entries except for the row that caused the failure. • SchemaViolation. When an operation fails, the data provider assigns this value to the row that caused the failure.
BindByName	<p>Specifies how the data provider processes named parameters when executing a stored procedure. The application can use named parameters or use default values for parameters to the stored procedure.</p> <p>If set to true, the data provider uses the names of parameters supplied in the <code>PsqlParameter</code> objects for the parameter bindings to the Zen server. See example for CommandText.</p> <p>Alternatively, the user can specify a default value for a named parameter using either of the following methods:</p> <ul style="list-style-type: none"> • The application binds the parameters using named parameters, but does not add a <code>PsqlParameter</code> object to the <code>PsqlParameterCollection</code> for the parameters for which the application wants to use the default value. • The application sets the <code>Value</code> property of the <code>PsqlParameter</code> object to null. The data provider does not send this parameter to the server and uses the parameter's default value when executing the stored procedure. <p>When <code>BindByName</code> is set to true and the <code>Parameter Mode</code> connection string option is defined as <code>BindByName</code> or <code>BindByOrdinal</code>, those values defined in the connection string are overridden for the lifetime of the <code>Command</code> object.</p> <p>If set to false (the initial default), the data provider ignores the names for the parameters supplied in the <code>PsqlParameter</code> objects, and assumes that the parameters are in the same order as they were specified in the <code>Create Procedure</code> statement.</p>

Property	Description
CommandText	<p>Gets or sets the text command to run against the data source.</p> <p>When using stored procedures, set CommandText to the name of the stored procedure, for example:</p> <pre>cmd.CommandType = System.Data.CommandType.StoredProcedure; cmd.CommandText = "call EnrollStudent(!!Stud_id!!,!!Class_Id!!,!!GPA!!)"; cmd.BindByName = true; PsqlParameter Class_Id = new PsqlParameter(); Class_Id.Value = 999; Class_Id.ParameterName = "!!Class_Id!!"; PsqlParameter Stud_id = new PsqlParameter(); Stud_id.Value = 1234567890; Stud_id.ParameterName = "!!Stud_id!!"; PsqlParameter GPA = new PsqlParameter(); GPA.Value = 3.2; GPA.ParameterName = "!!GPA!!"; cmd.Parameters.Add(Class_Id); cmd.Parameters.Add(Stud_id); cmd.Parameters.Add(GPA);</pre>
CommandTimeout	<p>Gets or sets the wait time before terminating the attempt to execute a command and generating an error.</p> <p>The initial default is 30 seconds.</p> <p>We recommend that the application sets the CommandTimeout property to a value that is greater than the largest default timeout value on the server. This ensures that the application gets a more meaningful reply in case of a timeout.</p>
CommandType	<p>Indicates or specifies how the CommandText property is interpreted.</p> <p>To use stored procedures, set CommandType to StoredProcedure.</p>
Connection	Gets or sets the IDbConnection used by this instance of the IDbCommand.
Parameters	Gets the PsqlParameterCollection.
RowsetSize	<p>Limits the number of rows returned by any query executed on this Command object to the value specified at execute time. The data type for the Read-Write property is signed integer.</p> <p>Valid values are 0 to 2147483647.</p> <p>If set to 0 (the initial default), the data provider does not limit the number of rows returned.</p>

Property	Description
StatementCacheBehavior	<p>Retrieves the statement cache behavior or sets the statement cache behavior to one of the values in the PsqlStatementCacheBehavior enumeration. See Enabling Statement Caching for more information.</p> <p>If set to Implicit (the default) and the Statement Cache Mode connection string option is set to Auto, statement caching occurs implicitly.</p> <p>If set to Cache and the Statement Cache Mode connection string option is set to ExplicitOnly, statements identified as Cache are cached.</p> <p>If set to DoNotCache, statement caching does not occur.</p>
Transaction	Gets or sets the transaction in which the PsqlCommand object executes.
UpdatedRowSource	<p>Gets or sets how command results are applied to the DataRow, when used by the Update method of a DataAdapter.</p> <p>When the UpdateBatchSize property is set to a value other than 1, the UpdatedRowSource property for UpdateCommand, DeleteCommand, and InsertCommand must be set to None or OutputParameters.</p> <p>If set to None, any returned parameters or rows are ignored.</p> <p>If set to OutputParameters, output parameters are mapped to the changed row in the DataSet.</p>

[Zen ADO.NET Entity Framework Data Providers](#) describes the public methods supported by the PsqlCommand object.

Method	Description
Cancel	Attempts to cancel the execution of an IDbCommand.
CreateParameter	Creates a new instance of an IDbDataParameter object.
Dispose	Releases the resources used by the component. Overloaded.
ExecuteNonQuery	Executes a SQL statement against the PsqlConnection object, and returns the number of rows affected. This method is intended for commands that do not return results.
ExecuteReader	Executes the CommandText against the connection and builds an IDataReader.
ExecuteScalar	Executes the query, and returns the first row of the resultset that the query returns. Any additional rows or columns are ignored.
Prepare	<p>Creates a prepared version of the command on an instance of Zen.</p> <p>Note: The Prepare method has no effect in this release of the data provider.</p>

PsqlCommandBuilder Object

Using a PsqlCommandBuilder object can have a negative effect on performance. Because of concurrency restrictions, the PsqlCommandBuilder can generate highly inefficient SQL statements. The end user can often write more efficient update and delete statements than those that the PsqlCommandBuilder generates.

The following table describes the public properties supported by the PsqlCommandBuilder object.

Property	Description
DataAdapter	Gets or sets the PsqlDataAdapter object associated with this PsqlCommandBuilder.

See [Type Mapping for Model First](#) for the public methods supported by the PsqlCommandBuilder object.

Method	Description
DeriveParameters	Populates the specified PsqlCommand object's Parameters collection with parameter information for a stored procedure specified in the PsqlCommand.
GetDeleteCommand	Gets the automatically generated PsqlCommand object required to perform deletions on the database when an application calls Delete on the PsqlDataAdapter.
GetInsertCommand	Gets the automatically generated PsqlCommand object required to perform inserts on the database when an application calls Insert on the PsqlDataAdapter.
GetUpdateCommand	Gets the automatically generated PsqlCommand object required to perform updates on the database when an application calls Update on the PsqlDataAdapter.
QuoteIdentifier	Given an unquoted identifier in the correct catalog case, returns the correct quoted form of that identifier, including properly escaping any embedded quotes in the identifier.
UnquoteIdentifier	Given a quoted identifier, returns the correct unquoted form of that identifier, including properly un-escaping any embedded quotes in the identifier.

PsqlConnection Object

The PsqlConnection object supports the public constructors described in the following table.

Property	Description
PsqlConnection()	Initializes a new instance of the PsqlConnection class.
PsqlConnection(string connectionString)	Initializes a new instance of the PsqlConnection class when given a string that contains the connection string.
PsqlConnection(string connectionString, PsqlCredential credential)	Initializes a new instance of the PsqlConnection class given a connection string and a PsqlCredential object that contains the user ID and password.

The PsqlConnection object supports the public properties described in the following table. Some properties return the values specified for the corresponding connection string option. Unlike the connection string options, the PsqlConnection property names do not include a space.

Property	Description
ConnectionString	Gets or sets the string used to open a database. See Table for a description of the values you can set.
ConnectionTimeout	<p>Gets the time to wait while trying to establish a connection before the data provider terminates the attempt and generates an error.</p> <p>You can set the amount of time a connection waits to time out by using the ConnectTimeout property or the Connection Timeout connection string option.</p> <p>If connection failover is enabled (the AlternateServers property defines one or more alternate database servers), this property applies to each attempt to connect to an alternate server. If connection retry is also enabled (the Connection Retry Count connection string option is set to an integer greater than 0), the ConnectionTimeout property applies to each retry attempt.</p>
Credential	Provides a more secure way to specify the password for connecting to a Zen server. PsqlCredential is comprised of a user ID and a password that will be used for connecting to a Zen server. The SecureString object which holds the password should be marked read-only.
Database	Gets the name of the current database or the database to be used when a connection is open.
Host	Returns the value specified for the Host connection string option. Read-only.
Port	Returns the value specified for the Port connection string option. Read-only.

Property	Description
ServerDSN	Returns the value specified for the Server DSN connection string option. Read-only.
ServerName	Returns the value specified for the Server Name connection string option. Read-only.
ServerVersion	Returns a string containing the version of the Zen server to which this object is currently connected. If the PsqlConnection object is not currently connected, the data provider generates an InvalidOperationException exception.
State	Gets the current state of the connection.
StatisticsEnabled	Enables statistics gathering. If set to True, enables statistics gathering for the current connection.

The following table describes the public methods of PsqlConnection.

Method	Description
BeginTransaction	Begins a database transaction. When using the overloaded BeginTransaction(IsolationLevel) method, the data provider supports isolation levels ReadCommitted and Serializable. See Isolation Levels for more information.
ChangeDatabase	Changes the current database for an open Connection object.
ClearAllPools	Empties the connection pools for the data provider.
ClearPool	Clears the connection pool that is associated with connection. If additional connections associated with the connection pool are in use at the time of the call, they are marked appropriately and are discarded when Close is called on them.
Close	Closes the connection to the database.
CreateCommand	Creates and returns a PsqlCommand object associated with the PsqlConnection.
Dispose	Releases the resources used by the PsqlConnection object.
Open	Opens a database connection with the settings specified by the ConnectionString property of the PsqlConnection object.
ResetStatistics	Resets all values to zero in the current statistics gathering session on the connection. When the connection is closed and returned to the connection pool, statistics gathering is switched off and the counts are reset.

Method	Description
RetrieveStatistics	Retrieves a set of statistics for a connection that is enabled for statistics gathering (see the StatisticsEnabled property). The set of name=value pairs returned forms a "snapshot in time" of the state of the connection when the method is called.

You can use the InfoMessage event of the PsqlConnection object to retrieve warnings and informational messages from the database. If the database returns an error, an exception is thrown. Clients that want to process warnings and informational messages sent by the database server should create a PsqlInfoMessageEventHandler delegate to register for this event.

The InfoMessage event receives an argument of type PsqlInfoMessageEventArgs containing data relevant to this event.

PsqlConnectionStringBuilder Object

PsqlConnectionStringBuilder property names are the same as the connection string option names of the PsqlConnection.ConnectionString property. However, the connection string option name can have spaces between the words. For example, the connection string option name Min Pool Size is equivalent to the property name MinPoolSize.

The basic format of a connection string includes a series of keyword/value pairs separated by semicolons. The following example shows the keywords and values for a simple connection string for the ADO.NET Data Provider:

```
"Server DSN=SERVERDEMO;Host=localhost"
```

Connection String Properties

The following table lists the properties that correspond to the connection string options supported by the Zen data providers, and describes each property.

Property	Description
AlternateServers	<p>Specifies a list of alternate database servers to which the data provider will try to connect if the primary database server is unavailable. Specifying a value for this connection string option enables connection failover for the data provider.</p> <p>The value you specify must be in the form of a string that defines connection information for each alternate server. You must specify the name or the IP address of each alternate server and the port number, if you are not using the default port value of 1583. The string has the format:</p> <pre>"Host=hostvalue;Port=portvalue[, ...]"</pre> <p>For example, the following Alternate Servers value defines two alternate servers for connection failover:</p> <pre>Alternate Servers="Host=AcctServer;Port=1584, Host=123.456.78.90;Port=1584"</pre> <p>See Using Connection Failover for a discussion of connection failover and information about other connection string options that you can set for this feature.</p>
ConnectionPoolBehavior	<p>{LeastRecentlyUsed MostRecentlyUsed LeastFrequentlyUsed MostFrequentlyUsed}. Specifies the order in which a connection is removed from the connection pool for reuse, based on how frequently or how recently the connection has been used.</p> <p>If set to MostRecentlyUsed, the data provider uses a Last In First Out (LIFO) approach to return the connection that was returned to the pool most recently.</p> <p>If set to LeastRecentlyUsed (the initial default), the data provider uses a First In First Out (FIFO) approach to return the connection that has been in the pool for the longest time. This value ensures a balanced use of connections in the pool.</p> <p>If set to MostFrequentlyUsed, the data provider returns the connection with the highest use count. This value enables applications to give preference to the most seasoned connection.</p> <p>If set to LeastFrequentlyUsed, the data provider returns the connection with the lowest use count. This value ensures a balanced use of connections in the pool.</p>

Property	Description
ConnectionReset	<p>{True False}. Specifies whether a connection that is removed from the connection pool for reuse by an application will have its state reset to the initial configuration settings of the connection. Resetting the state impacts performance because the new connection must issue additional commands to the server, for example, resetting the current database to the value specified at connect time.</p> <p>If set to False (the initial default), the data provider does not reset the state of the connection.</p>
ConnectionRetryCount	<p>Specifies the number of times the data provider tries to connect to the primary server, and, if specified, the alternate servers after the initial unsuccessful attempt.</p> <p>The value can be any integer from 0 to 65535.</p> <p>If set to 0 (the initial default), the data provider does not try to reconnect after the initial unsuccessful attempt.</p> <p>If a connection is not established during the retry attempts, the data provider returns an error that is generated by the last server to which it attempted to connect.</p> <p>This option and Connection Retry Delay, which specifies the wait interval between attempts, can be used in conjunction with connection failover. See Using Connection Failover for a discussion of connection failover and for information about other connection string options that you can set for this feature.</p>
ConnectionRetryDelay	<p>Specifies the number of seconds the data provider waits after the initial unsuccessful connection attempt before retrying a connection to the primary server, and, if specified, the alternate servers.</p> <p>The value can be any integer from 0 to 65535.</p> <p>The initial default is 3 (seconds). If set to 0, there is no delay between retrying the connection.</p> <p>Note: This option has no effect unless the Connection Retry Count connection string option is set to an integer value greater than 0.</p> <p>This option and the Connection Retry Count connection string option, which specifies the number of times the data provider attempts to connect after the initial attempt, can be used in conjunction with connection failover. See Using Connection Failover for a discussion of connection failover and for information about other connection string options that you can set for this feature.</p>

Property	Description
ConnectionTimeout	<p>Specifies the number of seconds after which the attempted connection to the server will fail if not yet connected. If connection failover is enabled, this option applies to each connection attempt.</p> <p>If set to 0, the data provider never times out on a connection attempt.</p> <p>The initial default is 15 seconds.</p>
DatabaseName	<p>Specifies the internal name of the database to which you want to connect. Use this option when you need to connect to a Zen data source for which a ServerDSN has not been defined.</p> <p>The default value is an empty string.</p> <p>Note: Do not combine the Database Name and Server DSN connection string options in a connection string.</p> <p>Alias: DBQ</p>
DbFileDirectoryPath	<p>Note: This option is supported only for Zen ADO.NET Entity Framework Core data providers.</p> <p>It determines in which directory on the database server the database files are created.</p> <p>The default value is an empty string.</p>
EnableIPv6	<p>Provides backward compatibility for connecting to a Zen server using an IPv4 address.</p> <p>If set to True, a client with IPv6 protocol installed can identify itself to the server using either an IPv4 address or an IPv6 address.</p> <p>If set to False, the clients runs in the backward compatibility mode. The client always identifies itself to the server using an IPv4 address.</p> <p>The default value for 4.0 is set to True.</p> <p>For more information about IPv6 formats, see IPv6 in <i>Getting Started with Zen</i>.</p>
EnableTrace	<p>{0 1}. Specifies whether tracing is enabled.</p> <p>If set to 0 (the initial default), tracing is not enabled.</p>
Encoding	<p>Specifies an IANA name or Windows code page number to be used for translating the string data stored in the database.</p> <p>The default value is an empty string; the current Windows Active Code Page (ACP) is used.</p>

Property	Description
Encrypt	<p>{If Needed Always Never}. Determines whether the data provider uses Encrypted Network Communications, also known as wire encryption.</p> <p>If set to Always, the data provider uses encryption, or, if the server does not allow wire encryption, returns an error.</p> <p>If set to Never, the data provider does not use encryption and returns an error if wire encryption is required by the server.</p> <p>If set to IfNeeded (the default), the data provider uses the default setting on the server.</p> <p>Note: This option may adversely affect performance because of the additional overhead, mainly CPU usage, required to encrypt and decrypt data.</p>
Encryption	<p>{Low Medium High}. Determines the minimum level of encryption allowed by the data provider.</p> <p>The initial default is Medium.</p> <p>The meaning of these values depends on the encryption module used. With the default encryption module, these values correspond to 40-, 56-, and 128-bit encryption, respectively.</p>
Enlist	<p>{True False}. Specifies whether the data provider automatically attempts to enlist the connection in creating the thread's current transaction context.</p> <p>Note: Because Zen does not support distributed transactions, any attempt to enlist the connection in the thread's current transaction context will fail.</p> <p>If set to False, the data provider does not automatically attempt to enlist the connection.</p> <p>If set to True (the initial default), the data provider returns an error on the connection if a current transaction context exists. If a current transaction context does not exist, the data provider raises a warning.</p>
Host	<p>Specifies the name or the IP address of the Zen database server to which you want to connect. For example, you can specify a server name such as Accountingserver. Or, you can specify an IP address such as 199.226.22.34 (IPv4) or 1234:5678:0000:0000:0000:0000:9abc:def0 (IPv6).</p> <p>The initial default value is an empty string.</p> <p>Alias: Server, Server Name</p>

Property	Description
InitialCommandTimeout	<p>Specifies the default wait time (timeout in seconds) before the data provider terminates the attempt to execute the command and generates an error. This option provides the same functionality as the PsqlCommand object's CommandTimeout property without the need to make changes to the application code. Subsequently, an application can use the CommandTimeout property to override the Initial Command Timeout connection string option.</p> <p>The initial default value is 30. If set to 0, the query never times out.</p> <p>For example, in the following C# code fragment, the connection string instructs the application to wait 60 seconds before terminating the attempt to execute the command. The application then specifies a CommandTimeout of 45 seconds, which overrides the value set in the connection string:</p> <pre> PsqlCommand command = new PsqlCommand(); PsqlConnection conn = new PsqlConnection("...; Initial Command Timeout=60; ..."); conn .Open(); command.Connection = connection; // command.CommandTimeout returns 60; command.CommandTimeout = 45; // command.CommandTimeout returns 45 command = new PsqlCommand(); command.CommandTimeout = 45; command.Connection = conn; // command.CommandTimeout still returns 45 </pre> <p>Note: Set the Initial Command Timeout option to a value that is greater than the largest default deadlock detection and timeout value on the server. This ensures that the application receives a more meaningful reply in case of a timeout.</p>
InitializationString	<p>Specifies one statement that will be issued immediately after connecting to the database to manage session settings.</p> <p>The initial default is an empty string.</p> <p>Example: To handle CHAR columns that are padded with NULLs, set the value to:</p> <pre> Initialization String=SET ANSI_PADDING ON </pre> <p>Note: If the statement fails to execute for any reason, the connection to the Zen server fails. The data provider throws an exception that contains the errors returned from the server.</p>

Property	Description
LoadBalanceTimeout	<p>Specifies the number of seconds to keep connections in a connection pool. The pool manager periodically checks all pools, and closes and removes any connection that exceeds its lifetime. The MinPoolSize option can cause some connections to ignore this value. See Removing Connections from a Pool for a discussion of connection lifetimes.</p> <p>The value can be any integer from 0 to 65 335.</p> <p>If set to 0, (the initial default), the connections have the maximum timeout.</p> <p>Alias: Connection Lifetime</p>
LoadBalancing	<p>{True False}. Determines whether the data provider uses client load balancing in its attempts to connect to primary and alternate database servers. The list of alternate servers is specified by the Alternate Servers connection option.</p> <p>If set to True, the data provider attempts to connect to the database servers in random order. See Using Client Load Balancing for more information about load balancing.</p> <p>If set to False (the initial default), client load balancing is not used and the data provider connects to each server based on its sequential order (primary server first, then, alternate servers in the order they are specified).</p> <p>Note: This option has no effect unless alternate servers are defined for the Alternate Servers connection string option.</p> <p>The Load Balancing connection string option is an optional setting that you can use in conjunction with connection failover. See Using Connection Failover for more information for a discussion of connection failover and for information about other connection options that you can set for this feature.</p>
MaxPoolSize	<p>Specifies the maximum number of connections within a single pool. When the maximum number is reached, no additional connections can be added to the connection pool. The Max Pool Size Behavior connection string option can cause some connections to ignore this value temporarily.</p> <p>The value can be any integer from 1 to 65335.</p> <p>The initial default is 100.</p>

Property	Description
MaxPoolSizeBehavior	<p>{SoftCap HardCap}. Specifies whether the data provider can exceed the number of connections specified by the Max Pool Size option when all connections in the connection pool are in use.</p> <p>If set to SoftCap, the number of connections created can exceed the value set for Max Pool Size, but the number of connections pooled does not. When the maximum connections for the pool are in use and a connection request is received, the data provider creates a new connection. If a connection is returned to a pool that is full and contains idle connections, the pooling mechanism selects a connection to be discarded so the connection pool never exceeds the Max Pool Size.</p> <p>If set to HardCap (the initial default), when the maximum number of connections allowed in the pool are in use, any new connection requests wait for an available connection until the Connection Timeout is reached.</p>
MaxStatementCacheSize	<p>Specifies the maximum number of statements that can be held in the statement cache. The value can be 0, or any integer greater than 1. Setting the cache size to 0 disables statement caching.</p> <p>The initial default is 10.</p> <p>In most cases, using statement caching results in improved performance. See the "Performance Considerations" topic for your data provider for information on how this option can affect performance.</p>
MinPoolSize	<p>Specifies the number of connections created when a connection pool is initialized and the minimum number of connections that will be kept in the pool. The connection pool retains this number of connections even when some connections have exceeded their LoadBalanceTimeout value.</p> <p>The value can be any integer from 0 to 65535.</p> <p>If set to 0 (the initial default), when the connection is closed and sent to the connection pool, the pool retains only the original connection used to create the pool.</p> <p>If set to an integer from 1 to 65535, the specified number of duplicates of the connection are placed in the pool.</p> <p>See the "Performance Considerations" topic for your data provider for information on how pooling can affect performance.</p>

Property	Description
ParameterMode	<p>Specifies the behavior of native parameter markers and binding. This allows applications to reuse provider-specific SQL code and simplifies migration to Zen ADO.NET data providers.</p> <p>If set to ANSI (the initial default), the ? character is processed as a parameter marker and bound as ordinal. Applications can toggle the behavior of the BindByName property on a per-command basis.</p> <p>If set to BindByOrdinal, native parameter markers are used and are bound as ordinal for stored procedures and standard commands.</p> <p>If set to BindByName, native parameter markers are used and are bound by name for stored procedures and standard commands.</p> <p>Note: This option is not supported for Zen ADO.NET Entity Framework data providers.</p>
Password	<p>Specifies a case-insensitive password used to connect to your Zen database. A password is required only if security is enabled on your database. If so, contact your system administrator to get your password.</p> <p>Alias: PWD</p>
PersistSecurityInfo	<p>{True False}. Specifies whether to display security information in clear text in the ConnectionString property.</p> <p>If set to True, the value of the Password connection string option is displayed in clear text.</p> <p>If set to False (the initial default), the data provider does not display the password in the connection string.</p>
Pooling	<p>{True False}. Specifies whether connections are pooled. See Using Connection Pooling for more information about connection pooling.</p> <p>If set to True (the initial default), connection pooling is enabled.</p> <p>See the "Performance Considerations" topic for your data provider for information on how pooling can affect performance.</p>
Port	<p>Specifies the TCP port of the listener running on the Zen database.</p> <p>The default port number is 1583.</p>
PVTranslate	<p>{Auto Nothing}. Specifies whether the client should negotiate a compatible encoding with the server.</p> <p>If set to Auto, the data provider will set the Encoding connection property to the database code page. In addition, SQL query text will be sent to the engine using UTF-8 encoding instead of the data encoding. This preserves NCHAR string literals in the query text.</p> <p>If set to Nothing (the default), the setting for the Encoding connection property is used.</p>

Property	Description
SchemaCollectionTimeout	<p>Specifies the number of seconds after which an attempted schema collection operation fails if it is not yet completed.</p> <p>If set to 0, the data provider never times out on a schema collection operation attempt.</p> <p>The initial default is 120.</p>
SchemaOptions	<p>Specifies additional database metadata that can be returned. By default, the data provider prevents the return of some available performance-expensive database metadata to optimize performance. If your application needs this database metadata, specify the name or hexadecimal value of the metadata.</p> <p>This option can affect performance.</p> <p>See Date and Time Canonical Functions for the name and hexadecimal value of the database metadata that the data provider can add.</p> <p>If set to ShowColumnDefaults or 0x04, column defaults are returned.</p> <p>If set to ShowParameterDefaults or 0x08, column defaults are returned.</p> <p>If set to FixParameterDirections or 0x10, procedure definitions are returned.</p> <p>If set to ShowProcedureDefinitions or 0x20, procedure definitions are returned.</p> <p>If set to ShowViewDefinitions or 0x40, view definitions are returned.</p> <p>If set to ShowAll or 0xFFFFFFFF (the initial default), all database metadata is returned.</p> <p>For example, to return descriptions of procedure definitions, specify Schema Options=ShowProcedureDefinitions or Schema Options=0x20.</p> <p>To show more than one piece of the omitted database metadata, specify either a comma-separated list of the names, or the sum of the hexadecimal values of the column collections that you want to restrict. For example, to return descriptions of procedure definitions and view definitions (hexadecimal values 0x20 and 0x40, respectively), specify Schema Options=ShowProcedureDefinitions, ShowViewDefinitions or Schema Options=0x60.</p> <p>Note: This connection string option may adversely affect performance. See documentation on performance considerations for your data provider for more information.</p>
ServerDSN	<p>Specifies the name of the data source on the server, such as Server DSN=SERVERDEMO.</p> <p>The default value is DEMODATA.</p> <p>Note: Do not combine the Database Name and Server DSN connection string options in a connection string.</p>

Property	Description
StatementCacheMode	<p>Specifies the statement cache mode. The statement cache mode controls the behavior of the statement cache. Statements can be cached automatically or only cached when a command is explicitly marked.</p> <p>If set to Auto, statement caching is enabled for statements marked as Implicit by the PsqlCommand property StatementCacheBehavior. These commands have a lower priority than that of explicitly marked commands, that is, if the statement pool reaches its maximum number of statements, the statements marked Implicit are removed from the statement pool first to make room for statements marked Cache.</p> <p>If set to ExplicitOnly (the initial default), only statements that are marked Cache by the StatementCacheBehavior property are cached.</p> <p>In most cases, enabling statement caching results in improved performance. See Performance Considerations for information on how this option can affect performance of the ADO.NET data provider.</p> <p>Note: This option is not supported for Zen ADO.NET Entity Framework data providers.</p>
Timestamp	<p>{DateTime String}. Specifies whether Zen time stamps are stored and retrieved as strings in the data provider.</p> <p>If set to DateTime or not defined (the default), the data provider maps time stamps to the .NET DateTime type. This setting may be appropriate when native precision is required, for example, when using the PsqlCommandBuilder with a time stamp.</p> <p>If set to String, the time stamps are returned as strings. The data provider maps Zen time stamps to the .NET String type.</p>
TimeType	<p>{DateTime TimeSpan}. Specifies whether Zen Times are retrieved as Timespan or DateTime in the ADO.NET data provider.</p> <p>If set to As DateTime, the data provider maps the SQL type TIME to the .NET type System.DateTime.</p> <p>If set to As TimeSpan, the data provider maps the SQL type TIME to the .NET type System.DateTimespan.</p>
TraceFile	<p>Specifies the path and file name of the trace file.</p> <p>The initial default is an empty string. If the specified file does not exist, the data provider creates it.</p>
UseCurrentSchema	<p>This connection string option is not supported. Setting it will cause the data provider to throw an exception.</p>
UserID	<p>Specifies the default Zen user name used to connect to your Zen database.</p> <p>Alias: UID</p>

The following table lists the name and the hexadecimal value of the column collection that the data provider will omit from the returned data. To specify multiple values, specify a comma-separated list of the names, or the sum of the hexadecimal values of the column collections that you want to return.

Name	Hex Value	Collection/Column
ShowColumnDefaults ¹	0x04	Columns/COLUMN_DEFAULT
ShowParameterDefaults	0x08	ProcedureColumns//PARAMETER_DEFAULT
FixParameterDirections	0x10	ProcedureColumns/PARAMETER_TYPE
ShowProcedureDefinitions	0x20	Procedures/PROCEDURE_DEFINITION
ShowViewDefinitions	0x40	Views/VIEW_DEFINITION
ShowAll	0x7F	All

1. COLUMN_HAS_DEFAULT is always reported with a value of null.

The `PsqlConnectionStringBuilder` object has no provider-specific methods. For information about the methods supported, refer to the data provider's online help and the Microsoft .NET Framework SDK documentation.

PsqlCredential Object

The `PsqlCredential` object provides a secure way to log in using Zen server authentication. `PsqlCredential` is comprised of a user ID and a password recognized by the Zen server.

The password in a `PsqlCredential` object is of type `SecureString`, unlike `Connection String` where the password is unsecure until the provider reads it and converts it to `SecureString`. The password is handled in a secure way without writing it to memory. The string that stores the password is cleaned after use.

Note: Use `PsqlCredential` only when the authentication method requires the user ID and password. Also, if you are using Kerberos or Client, you should not use `PsqlCredential`. Finally, the `Connection String` should not include the user ID and password when the `Credential` object is being used.

The following code snippet shows how the `PsqlCredential` class can be used. The method used to convert a string into a `SecureString` in this example is one of many possible methods.

```
PsqlConnection con = null;  
PsqlCredential lobjCredential = null;
```

```

string userId = "ABCD";
SecureString password = ConvertToSecureString("XYXYX");
private static SecureString ConvertToSecureString(string value)
{
    var securePassword = new SecureString();
    foreach (char c in value.ToCharArray())
        securePassword.AppendChar(c);
    securePassword.MakeReadOnly();
    return securePassword;
}
try
{
    lobjCredential = new PsqlCredential(userId, password);
    con = new PsqlConnection("Host=nc-xxx;Port=xxxx;Database Name=xxxx",
        lobjCredential);
    con.Open();
    Console.WriteLine("Connection Successfully Opened...");
    con.Close();
}
catch (Exception e)
{
    Console.Write(e.Message)
}
finally
{
    if (null != con)
    {
        con.Close();
        con = null;
    }
    if (null != lobjCredential)
    {
        lobjCredential = null;
    }
}
}

```

The following table lists the provider-specific implementation of the public properties of the PsqlCredential object.

Property	Description
User ID	Returns the user ID component of the PsqlCredential object. Uses String data type. NULL and empty are invalid values.
Password	Returns the password component of the PsqlCredential object. Uses SecureString data type. NULL is an invalid value.

If you use the PsqlCredential object while opening the connection and want to use the same pooled connection, you need to reference the same PsqlCredential object so that the same connection is fetched from the available connection pool.

If you create a new credential object for each connection, the driver treats them separately and puts them into different connection pools, even if the same user ID and password are used.

PsqlDataAdapter Object

The PsqlDataAdapter object uses PsqlCommand objects to execute SQL commands on the Zen database, to load the DataSet with data, and to reconcile the changed data in the DataSet to the database.

The following table describes the public properties of PsqlDataAdapter.

Property	Description
UpdateBatchSize	<p>Gets or sets a value that specifies the number of commands that can be executed in a batch.</p> <p>If your application uses disconnected DataSets and updates those DataSets, you can positively influence performance by setting this property to a value greater than 1. By default, the data provider attempts to use the largest batch size possible. However, this may not equate to optimal performance for your application. Set the value based on the number of rows you typically update in the DataSet. For example, if you are updating less than 50 rows, a suggested setting for this property is 25.</p> <p>If set to 0, the PsqlDataAdapter uses the largest batch size the data source can support. The UpdatedRowSource property for the InsertCommand, UpdateCommand, and DeleteCommand must be set to None or OutputParameters.</p> <p>If set to 1, batch updating is disabled.</p> <p>If set to a value greater than 1, the specified number of commands are executed in a batch. The UpdatedRowSource property for the InsertCommand, UpdateCommand, and DeleteCommand must be set to None or OutputParameters.</p>
DeleteCommand	Gets or sets a SQL statement for deleting records from the Zen data source.
InsertCommand	Gets or sets a SQL statement used to insert new records into the Zen database.
SelectCommand	Gets or sets a SQL statement used to select records in the Zen database.
UpdateCommand	Gets or sets a SQL statement used to update records in the data source.

PsqlDataReader Object

The PsqlDataReader object is a forward-only cursor that retrieves read-only records from a database. Performance is better than using PsqlDataAdapter, but the result set cannot be modified.

The following table describes the public properties of PsqldataReader.

Property	Description
Depth	Gets a value indicating the depth of nesting for the current row.
HasRows	Gets a value indicating whether the PsqldataReader contains one or more rows.
IsClosed	Gets a value indicating whether the data reader is closed.
RecordsAffected	Gets the number of rows that were changed, inserted, or deleted by execution of the SQL statement.

The following table describes some of the public methods of the PsqldataReader.

Method	Description
Close	Closes the DataReader. Always call the Close method when you finish using the DataReader object.
GetSchemaTable	Returns a DataTable that describes the column metadata of the PsqldataReader. See PsqlCredential Object for more information.
NextResult	Advances the data reader to the next result when reading the results of batch SQL statements.
Read	Advances the IDataReader to the next result.

PsqlError Object

The PsqlError object collects information relevant to errors and warnings generated by the Zen server.

The following table describes the public properties supported by PsqlError.

Property	Description
Message	Gets the error message text returned from the Zen server.
Number	Gets the error number returned from the Zen server.
SQLState	<p>Gets the string representation of the SQLState when an exception is thrown by the Zen data provider, or 0 if the exception is not applicable to the error. This property is read-only.</p> <p>Note: For all of the ADO.NET client error messages which do not have any SQLstate information, S1000 is used as the default SQLState.</p>

PsqlErrorCollection Object

The PsqlErrorCollection object is created by a PsqlException to contain all the errors generated by the Zen server.

The following table provides the public provider-specific properties supported for the PsqlErrorCollection object. For information about other properties and methods supported, refer to the data provider's online help and the Microsoft .NET Framework SDK documentation.

Property	Description
Count	Gets the number of PsqlError objects generated by the Zen server.

The PsqlErrorCollection object supports the public methods described in the following table.

Method	Description
CopyTo	Copies the PsqlError objects from the ErrorCollection to the specified array.
GetEnumerator	Returns the IEnumerator interface for a given array.

PsqlException Object

Provider-specific exceptions are derived directly from the System.Data interface. Only the public properties and methods, for example, the Message property, are directly available on the System.Exception object in a generic sense. The SQLState and Number properties are only accessible through provider-specific code or by using reflection.

ADO.NET 2.0 introduced a new property on the DbException class, Data. This property returns a collection of key-value pair tuples that provide additional user-defined information about an exception. The ADO.NET Data Provider gets a collection of key/value pairs such as SQLState, Number, and ErrorPosition.

The Psql.Data.SqlClient prefix is applied to each key, for example:

```
Psql.Data.SqlClient.Data["SQLState"] = 28000;
```

The properties described in the following table apply to the last error generated, if multiple errors exist. The application should check the Count property of the PsqlErrorCollection returned in the

Errors property of this object to determine whether multiple errors occurred. See [PsqlErrorCollection Object](#) for more information.

Property	Description
Errors	Gets or sets a PsqlErrorCollection of one or more PsqlError objects.
Message	Specifies the error message text that is returned from the Zen server.
Number	Gets or sets the number returned from the Zen server.
SQLState	Returns the string representation of the SQLState when an exception is thrown by the Zen data provider, or 0 if the exception is not applicable to the error. This property is read-only.

PsqlFactory Object

Provider Factory classes allow users to program to generic objects. Once instantiated from DbProviderFactory, the factory generates the proper type of concrete class.

The following table lists the static methods used to accommodate choosing the ADO.NET Data Provider and instantiating its DbProviderFactory.

Method	Description
CreateCommand	Returns a strongly typed DbCommand instance.
CreateCommandBuilder	Returns a strongly typed DbCommandBuilder instance.
CreateConnection	Returns a strongly typed DbConnection instance.
CreateConnectionStringBuilder	Returns a strongly typed DbConnectionString instance.
CreateDataAdapter	Returns a strongly typed DbDataAdapter instance.
CreateDataSourceEnumerator	Returns a strongly typed PsqlDataSourceEnumerator instance.
CreateParameter	Returns a strongly typed DbParameter instance.

PsqlInfoMessageEventArgs Object

The PsqlInfoMessageEventArgs object is passed as an input to the PsqlInfoMessageEventHandler and contains information relevant to a warning generated by the Zen server.

The following table describes the public properties for `PsqlInfoMessageEventArgs`.

Property	Description
Errors	Specifies a <code>PsqlErrorCollection</code> that contains a collection of warnings sent from the Zen server. See PsqlErrorCollection Object for more information.
Message	Returns the text of the last message returned from the Zen server. The application should check the Count property of the <code>PsqlErrorCollection</code> returned in the Errors property of this object to determine whether multiple warnings occurred.

PsqlParameter Object

The `PsqlParameter` object represents a parameter to a `PsqlCommand` object.

The following table describes the public properties for `PsqlParameter`.

Property	Description
ArrayBindStatus	Determines whether any values in the array of <code>PsqlParameterStatus</code> entries should be bound as null. The <code>PsqlParameterStatus</code> enumeration contains the entry <code>NullValue</code> . When this property is not set, then no values are null. The length of the array should match the amount specified by the <code>PsqlCommand</code> object's <code>ArrayBindCount</code> property (see PsqlCommand Object). The initial default is null.
DbType	Gets or sets the <code>DbType</code> of the parameter.
Direction	Gets or sets a value that indicates whether the parameter is input-only, output-only, bidirectional, or the return value parameter of a stored procedure.
IsNullable	Gets or sets a value that indicates whether the parameter accepts null values.
ParameterName	Gets or sets the name of the <code>PsqlParameter</code> object.
Precision	Gets or sets the maximum number of digits used to represent the <code>Value</code> property.
Scale	Gets or sets the number of decimal places to which the <code>Value</code> property is resolved.
Size	Gets or sets the maximum size, in bytes, of the data within the column.
SourceColumn	Gets or sets the name of the source column that is mapped to the <code>DataSet</code> and used for loading or returning the <code>Value</code> property.

Property	Description
SourceColumnNullMapping	Sets or gets a value that indicates whether the source column is nullable.
SourceVersion	Gets or sets the DataRowVersion to use when loading the Value property.
Value	<p>Gets or sets the value of the parameter.</p> <p>The initial default value is null.</p> <p>Note: When array binding is enabled (see the ArrayBindCount property of the PsqlCommand Object), this property is specified as an array of values. Each array's length must match the value of the ArrayBindCount property. When specifying the array's values for binary type columns, the data will actually be specified as byte[]. This is an array of arrays of bytes. The data provider anticipates a "jagged" array as such when using parameter array binding with parameters.</p> <p>If set to null for a stored procedure parameter, the data provider does not send the parameter to the server. Instead, the default value for the parameter is used when executing the stored procedure.</p>

PsqlParameterCollection Object

The PsqlParameterCollection object represents a collection of parameters relevant to a PsqlCommand, and includes their mappings to columns in a DataSet.

The following table describes the public properties for PsqlParameterCollection.

Property	Description
Count	Gets the number of PsqlParameter objects in the collection.
IsFixedSize	Gets a value that indicates whether the PsqlParameterCollection has a fixed size.
IsReadOnly	Gets a value that indicates whether the PsqlParameterCollection is read-only.
IsSynchronized	Gets a value that indicates whether the PsqlParameterCollection is thread-safe.
Item	Gets the parameter at the specified index. In C#, this property is the indexer for the IDataParameterCollection class.
SyncRoot	Gets the object used to synchronize access to the PsqlParameterCollection.

The following table provides the public methods for PsqlParameterCollection.

Method	Description
Contains	Gets a value that indicates whether a parameter in the collection has the specified source table name.
IndexOf	Gets the location of the IDataParameter within the collection.
RemoveAt	Removes the IDataParameter from the collection.

PsqlTrace Object

The PsqlTrace object is created by the application to debug problems during development. Setting the properties in the PsqlTrace object overrides the settings of the environment variables. For your final application, be sure to remove references to the PsqlTrace object.

The following code fragment creates a Trace object named MyTrace.txt. All subsequent calls to the data provider will be traced to that file.

```
PsqlTrace MyTraceObject = new PsqlTrace();
    MyTraceObject.TraceFile="C:\\MyTrace.txt";
    MyTraceObject.RecreateTrace = 1;
    MyTraceObject.EnableTrace = 1;
```

The following table describes the public properties for the PsqlTrace object.

Property	Description
EnableTrace	If set to 1 or higher, enables tracing. The initial default value is 0. Tracing is disabled.
RecreateTrace	If set to 1, recreates the trace file each time the application restarts. If set to 0 (the initial default), the trace file is appended
TraceFile	Specifies the path and name of the trace file. The initial default is an empty string. If the specified file does not exist, the data provider creates it.

Note: Setting EnableTrace starts the tracing process. Therefore, you must define the property values for the trace file before setting EnableTrace. Once the trace processing starts, the values of TraceFile and RecreateTrace cannot be changed.

The following table describes the public methods for PsqTrace.

Method	Description
DumpFootprints	Displays the footprint of all source files in a data provider.

PsqTransaction Object

The following table describes the public properties of PsqTransaction. The following table

Property	Description
Connection	Specifies the PsqConnection object associated with the transaction. See PsqConnection Object for more information.
IsolationLevel	Defines the isolation level for the entire transaction. If the value is changed, the new value is used at execution time.

describes the public methods of the PsqTransaction object.

Method	Description
Commit	When overridden in a derived class, returns the Exception that is the root cause of one or more subsequent exceptions.
Rollback	Cancels modifications made in a transaction before the transaction is committed.

Zen Common Assembly Classes

Zen ADO.NET data providers support additional classes that provide enhanced functionality, such as bulk load. All classes are created with 100% managed code. The following classes are provided in the Pervasive.Data.Common.dll assembly:

- [CsvDataReader](#)
- [CsvDataWriter](#)
- [DbBulkCopy](#)
- [DbBulkCopyColumnMapping](#)
- [DbBulkCopyColumnMapping](#)

The classes used for bulk loading implement the generic programming model. They can be used with any DataDirect Technologies ADO.NET data provider or ODBC driver that supports Zen Bulk Load and any supported database.

CsvDataReader

The CsvDataReader class provides the DataReader semantics for the CSV file format defined by Zen Bulk Load.

The following table lists the public properties for the CsvDataWriter object.

Property	Description
BulkConfigFile	<p>Specifies the CSV bulk configuration file that is produced when the WriteToFile method is called. A bulk load configuration file defines the names and data types of the columns in the bulk load data file in the same way as the table or result set from which the data was exported. A bulk load configuration file is supported by an underlying XML schema.</p> <p>The path may be fully qualified. Otherwise, the file is considered to exist in the current working directory.</p> <p>Note: This property can only be set prior to the Open() call and after the Close() call; otherwise, an InvalidOperationException is thrown.</p>

Property	Description
BulkFile	<p>Specifies the bulk load data file that contains the CSV-formatted bulk data. The file name is used for writing (exporting) and reading (importing) the bulk data. If the file name provided does not contain an extension, the ".csv" extension is assumed.</p> <p>The path may be fully qualified. Otherwise, the file is considered by default to exist in the current working directory. An <code>InvalidOperationException</code> is thrown if this value is not set.</p> <p>Note: This property can only be set prior to the <code>Open()</code> call and after the <code>Close()</code> call; otherwise, an <code>InvalidOperationException</code> is thrown.</p>
ReadBufferSize	<p>Specifies the size of the read buffer when using bulk load to import data from a data source.</p> <p>The initial default is 2048 KB.</p> <p>Values equal to or less than zero cause a <code>System.ArgumentOutOfRangeException</code> to be thrown.</p>
RowOffset	<p>Specifies the row from which to start the bulk load read. The <code>RowOffset</code> is relative to the first (1) row.</p> <p>The initial default is 1.</p> <p>Values equal to or less than zero cause a <code>System.ArgumentOutOfRangeException</code> to be thrown.</p> <p>Note: This property can only be set prior to the <code>Open()</code> call and after the <code>Close()</code> call; otherwise, an <code>InvalidOperationException</code> is thrown.</p>
SequentialAccess	<p>Determines whether columns are accessed in order for improved performance.</p> <p>The initial default is <code>False</code>.</p> <p>Note: This property can only be set prior to the <code>Open()</code> call and after the <code>Close()</code> call; otherwise, an <code>InvalidOperationException</code> is thrown.</p>

The following table lists the public methods for the `CsvDataReader` object.

Property	Description
Open	Opens the bulk file instance and associated metadatafile for processing.

CsvDataWriter

The `CsvDataWriter` class provides the `DataWriter` semantics of the CSV file format that is written by Zen Bulk Load.

For more information, refer to the data provider's online help.

The following table lists the public properties for the CsvDataWriter object.

Property	Description
BinaryThreshold	Specifies the threshold (in KB) at which separate files are generated to store binary data during a bulk unload. The Initial default is 32. Values less than zero throw a System.ArgumentOutOfRangeException to be thrown.
CharacterThreshold	Specifies the threshold (in KB) at which separate files are generated to store character data during a bulk unload. The initial default is 64. Values less than zero cause a System.ArgumentOutOfRangeException to be thrown.
CsvCharacterSetName	Specifies any of the supported IANA code page names that may be used as values. See IANA Code Page Mappings for the supported values. Applications can obtain the database character that was set using the SqlConnection.DatabaseCharacterSetName property. If an unrecognized CharacterSetName is used, an exception is thrown, declaring that invalid character set has been used. The initial default value is UTF-16. Note this property enforces the character set used in the CSV data file and overflow files added.

The following table lists the public methods for the CsvDataWriter object.

Property	Description
Open	Opens the bulk file instance and associated metadatafile for processing.
WriteToFile	Writes the contents of the IDataReader to the bulk data file.

DbBulkCopy

The DbBulkCopy class facilitates copying rows from one data source to another.

The DbBulkCopy object follows the de facto standard defined by the Microsoft SqlBulkCopy class, and has no provider-specific properties or methods. For information about the properties and methods supported, refer to the data provider's online help and the Microsoft .NET Framework SDK documentation.

DbBulkCopyColumnMapping

The `DbBulkCopyColumnMapping` class represents a column mapping from the data sources table to a destination table.

The `DbBulkCopyColumnMapping` object follows the de facto standard defined by the Microsoft `SqlBulkCopyColumnMapping` class, and has no provider-specific properties or methods. For information about the properties and methods supported, refer to the data provider's online help and the Microsoft .NET Framework SDK documentation.

DbBulkCopyColumnMappingCollection

The `DbBulkCopyColumnMappingCollection` class is a collection of `DbBulkCopyColumnMapping` objects.

The `DbBulkCopyColumnMappingCollection` object follows the de facto standard defined by the Microsoft `SqlBulkCopyColumnMappingCollection` class, and has no provider-specific properties or methods. For information about the properties and methods supported, refer to the data provider's online help and the Microsoft .NET Framework SDK documentation.

B. Getting Schema Information

Applications can request that data providers find and return metadata for a database. Schema collections specific to each data provider expose database schema elements such as tables and columns. The data provider uses the `GetSchema` method of the `Connection` class. You can also retrieve schema information from a result set, as described in [Columns Returned by the `GetSchemaTable` Method](#).

The data provider also includes provider-specific schema collections. Using the schema collection name `MetaDataCollections`, you can return a list of the supported schema collections, and the number of restrictions that they support.

Columns Returned by the GetSchemaTable Method

While a `PsqlDataReader` is open, you can retrieve schema information from the result set. The result set produced for `PsqlDataReader.GetSchemaTable()` returns the columns described in the following table, in the order shown.

Column	Description
ColumnName	Specifies the name of the column, which might not be unique. If the name cannot be determined, a null value is returned. This name reflects the most recent renaming of the column in the current view or command text.
ColumnOrdinal	Specifies the ordinal of the column, which cannot be null. The bookmark column of the row, if any, is 0. Other columns are numbered starting with 1.
ColumnSize	Specifies the maximum possible length of a value in the column. For columns that use a fixed-length data type, this is the size of the data type.
NumericPrecision	Specifies the precision of the column, which depends on how the column is defined in <code>ProviderType</code> . If <code>ProviderType</code> is a numeric data type, this is the maximum precision of the column. If <code>ProviderType</code> is not a numeric data type, the value is null.
NumericScale	Specifies the number of digits to the right of the decimal point if <code>ProviderType</code> is <code>DBTYPE_DECIMAL</code> or <code>DBTYPE_NUMERIC</code> . Otherwise, this is a null value. The value depends on how the column is defined in <code>ProviderType</code> .
DataType	Maps to the .NET Framework type of the column.
ProviderType	Specifies the indicator of the column data type. This column cannot contain a null value. If the data type of the column varies from row to row, this must be <code>Object</code> .
IsLong	Set if the column contains a BLOB that contains very long data. The setting of this flag corresponds to the value of the <code>IS_LONG</code> column in the <code>PROVIDER_TYPES</code> rowset for the data type. The definition of very long data is provider-specific.
AllowDBNull	Set if the consumer can set the column to a null value, or if the data provider cannot determine whether the consumer can set the column to a null value. Otherwise, no value is set. A column can contain null values, even if it cannot be set to a null value.

Column	Description
IsReadOnly	Determines whether a column can be changed. The value is true if the column can be modified; otherwise, the value is false.
IsRowVersion	Is set if the column contains a persistent row identifier that cannot be written to, and has no meaningful value except to identify the row.
IsUnique	Specifies whether the column constitutes a key by itself or if there is a constraint of type UNIQUE that applies only to this column. If set to true, no two rows in the base table (the table returned in BaseTableName) can have the same value in this column. If set to false (the initial default), the column can contain duplicate values in the base table.
IsKey	Specifies whether a set of columns uniquely identifies a row in the rowset. This set of columns may be generated from a base table primary key, a unique constraint, or a unique index. The value is true if the column is one of a set of columns in the rowset that, taken together, uniquely identify the row. The value is false if the column is not required to uniquely identify the row.
IsAutoIncrement	Specifies whether the column assigns values to new rows in fixed increments. If set to VARIANT_TRUE, the column assigns values to new rows in fixed increments. If set to VARIANT_FALSE (the initial default), the column does not assign values to new rows in fixed increments.
BaseSchemaName	Specifies the name of the schema in the database that contains the column. The value is null if the base schema name cannot be determined. The initial default is null.
BaseCatalogName	Specifies the name of the catalog in the data store that contains the column. A null value is used if the base catalog name cannot be determined. The initial default is null.
BaseTableName	Specifies the name of the table or view in the data store that contains the column. A null value is used if the base table name cannot be determined. The initial default is null.
BaseColumnName	Specifies the name of the column in the data store. This might be different than the column name returned in the ColumnName column if an alias was used. A null value is used if the base column name cannot be determined or if the rowset column is derived from, but is not identical to, a column in the database. The initial default is null.

Column	Description
IsAliased	Specifies whether the name of the column is an alias. The value true is returned if the column name is an alias; otherwise, false is returned.
IsExpression	Specifies whether the name of the column is an expression. The value true is returned if the column is an expression; otherwise, false is returned.
IsIdentity	Specifies whether the name of the column is an identity column. The value true is returned if the column is an identity column; otherwise, false is returned.
IsHidden	Specifies whether the name of the column is hidden. The value true is returned if the column is hidden; otherwise, false is returned.

Retrieving Schema Metadata with the GetSchema Method

Applications use the GetSchema method of the Connection object to retrieve Schema Metadata about a data provider and/or data source. Each provider implements a number of Schema collections, including the five standard metadata collections.

- [MetaDataCollections Schema Collections](#)
- [DataSourceInformation Schema Collection](#)
- [DataTypes Collection](#)
- [ReservedWords Collection](#)
- [Restrictions Collection](#)

Additional collections are specified and must be supported to return Schema information from the data provider.

See [Additional Schema Collections](#) for details about the other collections supported by the data providers.

Note: Refer to the .NET Framework documentation for additional background functional requirements, including the required data type for each ColumnName.

MetaDataCollections Schema Collections

The MetaDataCollections schema collection is a list of the schema collections available to the logged in user. The MetaDataCollection can return the supported columns described in the following table, in any order.

ColumnName	Description
CollectionName	The name of the collection to pass to the GetSchema method to return the collection.
NumberOfRestrictions	The number of restrictions that may be specified for the collection.
NumberOfIdentifierParts	The number of parts in the composite identifier/data base object name.

DataSourceInformation Schema Collection

The DataSourceInformation schema collection can return the supported columns, described in the following table, in any order. Note that only one row is returned.

ColumnName	Description
CompositeIdentifierSeparatorPattern	The regular expression to match the composite separators in a composite identifier.
DataSourceProductName	The name of the product accessed by the data provider.
DataSourceProductVersion	Indicates the version of the product accessed by the data provider, in the data source's native format.
DataSourceProductVersionNormalized	A normalized version for the data source. This allows the version to be compared with String.Compare().
GroupByBehavior	Specifies the relationship between the columns in a GROUP BY clause and the non-aggregated columns in the select list.
Host	The host to which the data provider is connected.
IdentifierCase	Indicates whether non-quoted identifiers are treated as case sensitive.
IdentifierPattern	A regular expression that matches an identifier and has a match value of the identifier.
OrderByColumnsInSelect	Specifies whether columns in an ORDER BY clause must be in the select list. A value of true indicates that they are required to be in the Select list, a value of false indicates that they are not required to be in the Select list.
ParameterMarkerFormat	A format string that represents how to format a parameter.
ParameterMarkerPattern	A regular expression that matches a parameter marker. It will have a match value of the parameter name, if any.
ParameterNameMaxLength	The maximum length of a parameter name in characters.
ParameterNamePattern	A regular expression that matches the valid parameter names.
QuotedIdentifierCase	Indicates whether quoted identifiers are treated as case sensitive.
QuotedIdentifierPattern	A regular expression that matches a quoted identifier and has a match value of the identifier itself without the quotation marks.

ColumnName	Description
StatementSeparatorPattern	A regular expression that matches the statement separator.
StringLiteralPattern	A regular expression that matches a string literal and has a match value of the literal itself.
SupportedJoinOperators	Specifies the types of SQL join statements that are supported by the data source.

DataTypes Collection

The following table describes the supported columns of the DataTypes schema collection. The columns can be returned in any order.

ColumnName	Description
ColumnSize	The length of a non-numeric column or parameter refers to either the maximum or the length defined for this type by the data provider.
CreateFormat	Format string that represents how to add this column to a data definition statement, such as CREATE TABLE.
CreateParameters	<p>The creation parameters that must be specified when creating a column of this data type. Each creation parameter is listed in the string, separated by a comma in the order they are to be supplied.</p> <p>For example, the SQL data type DECIMAL needs a precision and a scale. In this case, the creation parameters should contain the string "precision, scale".</p> <p>In a text command to create a DECIMAL column with a precision of 10 and a scale of 2, the value of the CreateFormat column might be DECIMAL({0},{1})" and the complete type specification would be DECIMAL(10,2).</p>
DataType	The name of the .NET Framework type of the data type.
IsAutoIncrementable	<p>Specifies whether values of a data type are auto-incremented.</p> <p>true: Values of this data type may be auto-incremented.</p> <p>false: Values of this data type may not be auto-incremented.</p>
IsBestMatch	<p>Specifies whether the data type is the best match between all data types in the data store and the .NET Framework data type indicated by the value in the DataType column.</p> <p>true: The data type is the best match.</p> <p>false: The data type is not the best match.</p>

ColumnName	Description
IsCaseSensitive	Specifies whether the data type is both a character type and case-sensitive. true: The data type is a character type and is case-sensitive. false: The data type is not a character type or is not case-sensitive.
IsConcurrencyType	true: The data type is updated by the database every time the row is changed and the value of the column is different from all previous values. false: The data type is not updated by the database every time the row is changed.
IsFixedLength	true: Columns of this data type created by the data definition language (DDL) will be of fixed length. false: Columns of this data type created by the DDL will be of variable length.
IsFixedPrecisionScale	true: The data type has a fixed precision and scale. false: The data type does not have a fixed precision and scale.
IsLiteralsSupported	true: The data type can be expressed as a literal. false: The data type cannot be expressed as a literal.
IsLong	true: The data type contains very long data. The definition of very long data is provider-specific. false: The data type does not contain very long data.
IsNullable	true: The data type is nullable. false: The data type is not nullable.
IsSearchable	true: The data type contains very long data. The definition of very long data is provider-specific. false: The data type does not contain very long data.
IsSearchableWithLike	true: The data type can be used with the LIKE predicate. false: The data type cannot be used with the LIKE predicate.
IsUnsigned	true: The data type is unsigned. false: The data type is signed.
LiteralPrefix	The prefix applied to a given literal.
LiteralSuffix	The suffix applied to a given literal.
MaximumScale	If the type indicator is a numeric type, this is the maximum number of digits allowed to the right of the decimal point. Otherwise, this is DBNull.Value.

ColumnName	Description
MinimumScale	If the type indicator is a numeric type, this is the minimum number of digits allowed to the right of the decimal point. Otherwise, this is <code>DBNull.Value</code> .
ProviderDbType	The provider-specific type value that should be used when specifying a parameter's type.
TypeName	The provider-specific data type name.

ReservedWords Collection

This schema collection exposes information about the words that are reserved by the database to which the data provider is connected. The following table describes the columns that the data provider supports.

ColumnName	Description
Reserved Word	Provider-specific reserved words.

Restrictions Collection

The Restrictions schema collection exposes information about the restrictions supported by the data provider that is currently connected to the database. The following table describes the columns returned by the data providers. The columns can be returned in any order.

The ADO.NET Data Provider uses standardized names for restrictions. If the data provider supports a restriction for a Schema method, it always uses the same name for the restriction.

The case sensitivity of any restriction value is determined by the underlying database, and can be determined by the `IdentifierCase` and `QuotedIdentifierCase` values in the [DataSourceInformation](#) collection (see [DataSourceInformation Schema Collection](#)).

ColumnName	Description
CollectionName	The name of the collection to which the specified restrictions apply.
RestrictionName	The name of the restriction in the collection.
RestrictionDefault	Ignored.
RestrictionNumber	The actual location in the collection restrictions for this restriction.
IsRequired	Specifies whether the restriction is required.

See [Additional Schema Collections](#) for the restrictions that apply to the each of the additional supported schema collections.

Additional Schema Collections

The Zen ADO.NET Data Provider supports the following additional schema collections:

- [Columns Schema Collection](#)
- [ForeignKeys Schema Collection](#)
- [Indexes Schema Collection](#)
- [PrimaryKeys Schema Collection](#)
- [ProcedureParameters Schema Collection](#)
- [Procedures Schema Collection](#)
- [TablePrivileges Schema Collection](#)
- [Tables Schema Collection](#)
- [Views Schema Collection](#)

Columns Schema Collection

Description: The Columns schema collection identifies the columns of tables (including views) defined in the catalog that are accessible to a given user. The following table identifies the columns of tables defined in the catalog that are accessible to a given user.

Number of restrictions: 3

Restrictions available: TABLE_CATALOG, TABLE_NAME, COLUMN_NAME

Sort order: TABLE_CATALOG, TABLE_NAME, ORDINAL_POSITION

Column Name	.NET Framework DataType ¹	Description
CHARACTER_MAXIMUM_LENGTH	Int32	<p>The maximum possible length of a value in the column. For character, binary, or bit columns, this is one of the following:</p> <ul style="list-style-type: none"> The maximum length of the column in characters, bytes, or bits, respectively, if one is defined. The maximum length of the data type in characters, bytes, or bits, respectively, if the column does not have a defined length. Zero (0) if neither the column or the data type has a defined maximum length, or if the column is not a character, binary, or bit column.
CHARACTER_OCTET_LENGTH	Int32	<p>The maximum length in octets (bytes) of the column, if the type of the column is character or binary.</p> <p>A value of zero (0) means the column has no maximum length or that the column is not a character or binary column.</p>
COLUMN_DEFAULT	String	The default value of the column.
COLUMN_HASDEFAULT	Boolean	<p>TRUE: The column has a default value.</p> <p>FALSE: The column does not have a default value, or it is unknown whether the column has a default value.</p>
COLUMN_NAME	String	The name of the column; this might not be unique.
DATA_TYPE	Object	<p>The indicator of the column data type.</p> <p>This value cannot be null.</p>
IS_NULLABLE	Boolean	<p>TRUE: The column might be nullable.</p> <p>FALSE: The column is known not to be nullable.</p>
NATIVE_DATA_TYPE	String	<p>The data source description of the type.</p> <p>This value cannot be null.</p>

Column Name	.NET Framework DataType ¹	Description
NUMERIC_PRECISION	Int32	If the column data type is of numeric data, this is the maximum precision of the column.
NUMERIC_PRECISION_RADIX	Int32	The radix indicates in which base the values in NUMERIC_PRECISION and NUMERIC_SCALE are expressed. It is only useful to return either 2 or 10.
NUMERIC_SCALE	Int16	If the column type is a numeric type that has a scale, this is the number of digits to the right of the decimal point.
ORDINAL_POSITION	Int32	The ordinal of the column. Columns are numbered starting from one.
PROVIDER_DEFINED_TYPE	Int32	The data source defined type of the column is mapped to the type enumeration of the data provider, for example, the PsqlDbType enumeration. This value cannot be null.
PROVIDER_GENERIC_TYPE	Int32	The provider-defined type of the column is mapped to the System.Data.DbType enumeration. This value cannot be null.
TABLE_CATALOG	String	The database name.
TABLE_NAME	String	The table name.
TABLE_OWNER	String	The table owner.

1. All classes are System.XXX. For example, System.String.

ForeignKeys Schema Collection

Description: The ForeignKeys schema collection identifies the foreign key columns defined in the catalog by a given user.

Number of restrictions: 2

Restrictions available: FK_TABLE_CATALOG, PK_TABLE_NAME

Sort order: FK_TABLE_CATALOG, FK_TABLE_NAME

Column Name	.NET Framework Datatype ¹	Description
DEFERRABILITY	String	The deferrability of the foreign key. The value is one of the following: <ul style="list-style-type: none">INITIALLY DEFERREDINITIALLY IMMEDIATENOT DEFERRABLE
DELETE_RULE	String	If a delete rule was specified, the value is one of the following: CASCADE: A referential action of CASCADE was specified. SET NULL: A referential action of SET NULL was specified. SET DEFAULT: A referential action of SET DEFAULT was specified. NO ACTION: A referential action of NO ACTION was specified.
FK_COLUMN_NAME	String	The foreign key column name.
FK_NAME	String	The foreign key name. This is a required restriction.
FK_TABLE_CATALOG	String	The catalog name in which the foreign key table is defined.
FK_TABLE_NAME	String	The foreign key table name. This is a required restriction.
FK_TABLE_OWNER	String	The foreign key table owner. This is a required restriction.
ORDINAL	Int32	The order of the column names in the key. For example, a table might contain several foreign key references to another table. The ordinal starts over for each reference; for example, two references to a three-column key would return 1, 2, 3, 1, 2, 3.
PK_COLUMN_NAME	String	The primary key column name.
PK_NAME	String	The primary key name.

Column Name	.NET Framework Datatype ¹	Description
PK_TABLE_CATALOG	String	The catalog name in which the primary key table is defined.
PK_TABLE_NAME	String	The primary key table name.
PK_TABLE_OWNER	String	The primary key table owner. This is a required restriction.
UPDATE_RULE	String	If an update rule was specified, one of the following: CASCADE: A referential action of CASCADE was specified. SET NULL: A referential action of SET NULL was specified. SET DEFAULT: A referential action of SET DEFAULT was specified. NO ACTION: A referential action of NO ACTION was specified.

1. All classes are System.XXX. For example, System.String

Indexes Schema Collection

Description: The Indexes schema collection identifies the indexes defined in the catalog that are owned by a given user.

Number of restrictions: 2

Restrictions available: TABLE_CATALOG, TABLE_NAME

Sort order: UNIQUE, TYPE, INDEX_CATALOG, INDEX_NAME, ORDINAL_POSITION

Column Name	.NET Framework DataType ¹	Description
CARDINALITY	Int32	The number of unique values in the index.
COLLATION	String	This is one of the following: ASC: The sort sequence for the column is ascending. DESC: The sort sequence for the column is descending.

Column Name	.NET Framework DataType ¹	Description
COLUMN_NAME	String	The column name.
FILTER_CONDITION	String	The WHERE clause that identifies the filtering restriction.
INDEX_CATALOG	String	The catalog name.
INDEX_NAME	String	The index name.
ORDINAL_POSITION	Int32	The ordinal position of the column in the index, starting with 1.
PAGES	Int32	The number of pages used to store the index.
TABLE_CATALOG	String	The catalog name.
TABLE_NAME	String	The table name.
TABLE_OWNER	String	The table owner.
TABLE_QUALIFIER	String	The table qualifier.
TYPE	String	The type of the index. One of the following values: <ul style="list-style-type: none"> • BTREE: The index is a B+-tree. • HASH: The index is a hash file using, for example, linear or extensible hashing. • CONTENT: The index is a content index. • OTHER: The index is some other type of index.
UNIQUE	Boolean	

1. All classes are System.XXX. For example, System.String.

PrimaryKeys Schema Collection

Description: The PrimaryKeys schema collection identifies the primary key columns defined in the catalog by a given user.

Number of restrictions: 2

Restrictions available: TABLE_CATALOG, TABLE_NAME

Sort order: TABLE_CATALOG, TABLE_NAME

Column Name	.NET Framework DataType ¹	Description
COLUMN_NAME	String	The primary key column name.
ORDINAL	Int32	The order of the column names in the key.
PK_NAME	String	The primary key name.
TABLE_CATALOG	String	The database name in which the table is defined.
TABLE_NAME	String	The table name.
TABLE_OWNER	String	The table owner.

1. All classes are System.XXX. For example, System.String.

ProcedureParameters Schema Collection

Description: The ProcedureParameters schema collection returns information about the parameters and return codes of procedures that are part of the Procedures collection.

Number of restrictions: 3

Restrictions available: PROCEDURE_CATALOG, PROCEDURE_NAME, PARAMETER_NAME

Sort order: PROCEDURE_CATALOG, PROCEDURE_NAME, ORDINAL_POSITION

Column Name	.NET Framework DataType ¹	Description
CHARACTER_MAXIMUM_LENGTH	Int32	The maximum length of the parameter.
CHARACTER_OCTET_LENGTH	Int32	The maximum length in octets (bytes) of the parameter, if the type of the parameter is character or binary. If the parameter has no maximum length, the value is zero (0). For all other types of parameters, the value is -1.

Column Name	.NET Framework DataType ¹	Description
DATA_TYPE	Object	The indicator of the column data type. This value cannot be null.
DESCRIPTION	String	The description of the parameter. For example, the description of the Name parameter in a procedure that adds a new employee might be Employee name.
IS_NULLABLE	Boolean	TRUE: The parameter might be nullable. FALSE: The parameter is not nullable.
NATIVE_DATA_TYPE	String	The data source description of the type. This value cannot be null.
NULLABLE	String	Denotes whether null value can be specified for the parameter. YES and NO are the two possible values.
NUMERIC_PRECISION	Int32	If the column data type is numeric, this is the maximum precision of the column. If the column data type is not numeric, this is DbNull.
NUMERIC_PRECISION_RADIX	Int32	Applicable when the column data type is numeric. The radix indicates in which base the values in NUMERIC_PRECISION and NUMERIC_SCALE are expressed. It is only useful to return either 2 or 10.
NUMERIC_SCALE	Int16	If the column data type is a numeric type that has a scale, this is the number of digits to the right of the decimal point. Otherwise, this is DbNull.
ORDINAL_POSITION	Int32	If the parameter is an input, input/output, or output parameter, this is the one-based ordinal position of the parameter in the procedure call. If the parameter is the return value, this is DbNull.

Column Name	.NET Framework DataType ¹	Description
PARAMETER_DEFAULT	String	<p>The default value of parameter.</p> <p>If the default value is a NULL, then the PARAMETER_HASDEFAULT column returns TRUE and the PARAMETER_DEFAULT column will not exist.</p> <p>If PARAMETER_HASDEFAULT is set to FALSE, then the PARAMETER_DEFAULT column will not exist.</p>
PARAMETER_HASDEFAULT	Boolean	<p>TRUE: The parameter has a default value.</p> <p>FALSE: The parameter does not have a default value, or it is unknown whether the parameter has a default value.</p>
PARAMETER_NAME	String	The parameter name. If the parameter is not named, this is DbNull.
PARAMETER_TYPE	String	<p>This is one of the following:</p> <p>INPUT: The parameter is an input parameter.</p> <p>INPUTOUTPUT: The parameter is an input/output parameter.</p> <p>OUTPUT: The parameter is an output parameter.</p> <p>RETURNVALUE: The parameter is a procedure return value.</p> <p>UNKNOWN: The parameter type is unknown to the data provider.</p>
PROCEDURE_CATALOG	String	The catalog name.
PROCEDURE_NAME	String	The procedure name.
PROCEDURE_COLUMN_NAME	String	The procedure column name.

Column Name	.NET Framework DataType ¹	Description
PROVIDER_DEFINED_TYPE	Int32	The data source defined type of the column as mapped to the type enumeration of the data provider, for example, the PSQldbType enumeration. This value cannot be null.
PROVIDER_GENERIC_TYPE	Int32	The data source defined type of the column as mapped to the System.Data.DbType enumeration. This value cannot be null.
SQL_DATETIME_SUB	Object	Applicable when the column data type is DateTime.

1. All classes are System.XXX. For example, System.String.

Procedures Schema Collection

Description: The Procedures schema collection identifies the procedures defined in the catalog. When possible, only procedures for which the connected user has execute permission should be returned.

Number of restrictions: 2

Restrictions available: PROCEDURE_CATALOG, PROCEDURE_NAME, PROCEDURE_TYPE

Sort order: PROCEDURE_CATALOG, PROCEDURE_NAME

Column Name	.NET Framework DataType ¹	Description
PROCEDURE_CATALOG	String	The database name.
PROCEDURE_NAME	String	The procedure name.
PROCEDURE_OWNER	String	The procedure owner.

Column Name	.NET Framework DataType ¹	Description
PROCEDURE_TYPE	String	This is one of the following: UNKNOWN: It is not known whether a value is returned. PROCEDURE: Procedure; no value is returned. FUNCTION: Function; a value is returned.

1. All classes are System.XXX. For example, System.String.

TablePrivileges Schema Collection

Description: The TablePrivileges schema collection identifies the privileges on tables defined in the catalog that are available to or granted by a given user.

Number of restrictions: 3

Restrictions available: TABLE_CATALOG, TABLE_NAME, GRANTEE

Sort order: TABLE_CATALOG, TABLE_NAME, PRIVILEGE_TYPE

Column Name	Type Indicator ¹	Description
GRANTEE	String	The user name (or PUBLIC) to whom the privilege has been granted.
PRIVILEGE_TYPE	String	The privilege type. This is one of the following types: <ul style="list-style-type: none"> • DELETE • INSERT • REFERENCES • SELECT • UPDATE
TABLE_CATALOG	String	The name of the database in which the table is defined.
TABLE_NAME	String	The table name.
TABLE_OWNER	String	The table owner.

1. All classes are System.XXX. For example, System.String.

Tables Schema Collection

Description: The Tables schema collection identifies the tables (including views) defined in the catalog that are accessible to a given user.

Number of Restrictions: 3

Restrictions Available: TABLE_CATALOG, TABLE_NAME, TABLE_TYPE

Sort order: TABLE_TYPE, TABLE_CATALOG, TABLE_NAME

Column Name	.NET Framework DataType ¹	Description
DESCRIPTION	String	A description of the table. If no description is associated with the column, the data provider returns <code>DBNull</code> .
TABLE_CATALOG	String	The name of the database in which the table is defined.
TABLE_NAME	String	The table name.
TABLE_OWNER	String	The table owner.
TABLE_TYPE	String	The table type. One of the following: <ul style="list-style-type: none">• ALIAS• GLOBAL TEMPORARY• LOCAL TEMPORARY• SYNONYM• SYSTEM TABLE• SYSTEM VIEW• TABLE• VIEW This column cannot contain an empty string.

1. All classes are System.XXX. For example, System.String.

Views Schema Collection

Description: The Views schema collection identifies the views defined in the catalog that are accessible to a given user.

Number of restrictions: 2

Restrictions available: TABLE_CATALOG, TABLE_NAME

Sort order: TABLE_CATALOG, TABLE_NAME

Column Name	Type Indicator ¹	Description
TABLE_CATALOG	String	The name of the database in which the table is defined.
TABLE_NAME	String	The table name.
TABLE_OWNER	String	The table owner.
TABLE_QUALIFIER	String	The table qualifier.
VIEW_DEFINITION	String	The view definition. This is a query expression.

1. All classes are System.XXX. For example, System.String.

C. SQL Escape Sequences for .NET

A number of language features, such as outer joins and scalar function calls, are commonly implemented by database management systems. The syntax for these features is often DBMS-specific, even when a standard syntax has been defined. The .NET support for escape sequences contain standard syntaxes for the following language features:

- Date, time, and timestamp literals
- Scalar functions such as numeric, string, and data type conversion functions
- Outer joins

The escape sequence used by .NET is:

```
{extension}
```

The escape sequence is recognized and parsed by the ADO.NET Data Provider, which replaces the escape sequences with data store-specific grammar.

Date, Time, and Timestamp Escape Sequences

The escape sequence for date, time, and timestamp literals is:

```
{literal-type 'value'}
```

where *literal-type* is one of the following:

literal-type	Description	Value Format
d	Date	yyyy-mm-dd
t	Time	hh:mm:ss [1]
ts	Timestamp	yyyy-mm-dd hh:mm:ss[.f...]

Note: If you receive an error while running a query in Visual Studio to insert data into the Date field of a table, ensure that your system's Date format is set to yyyy-mm-dd. If not, change it to yyyy-mm-dd.

Example

```
UPDATE Orders SET OpenDate={d '1997-01-29'}  
WHERE OrderID=1023
```

Scalar Functions

You can use scalar functions in SQL statements with the following syntax:

```
{fn scalar-function}
```

where *scalar-function* is a scalar function supported by the ADO.NET Data Provider.

Example

```
SELECT {fn UCASE(NAME)} FROM EMP
```

The following table lists the scalar functions supported.

String Functions	Numeric Functions	Timedate Functions	System Functions
ASCII	ABS	CURDATE	DATABASE
BIT_LENGTH	ACOS	CURRENT_DATE	USER
CHAR	ASIN	CURTIME	
CHAR_LENGTH	ATAN	CURRENT_TIME	
CHARACTER_LENGTH	ATAN2	CURRENT_TIMESTAMP	
CONCAT	CEILING	DAYNAME	
LCASE or LOWER	COS	DAYOFMONTH	
LEFT	COT	DAYOFYEAR	
LENGTH	DEGREES	EXTRACT	
LOCATE	EXP	HOUR	
LTRIM	FLOOR	MINUTE	
OCTET_LENGTH	LOG	MONTH	
POSITION	LOG10	MONTHNAME	
REPLACE	MOD	NOW	
REPLICATE	PI	QUARTER	
RIGHT	POWER	SECOND	
RTRIM	RADIANS	TIMESTAMPADD	
SPACE	RAND	TIMESTAMPDIFF	
STUFF	ROUND	WEEK	
SUBSTRING	SIGN	YEAR	
UCASE or UPPER	SIN		
	SQRT		
	TAN		
	TRUNCATE		

Outer Join Escape Sequences

.NET supports the SQL92 left, right, and full outer join syntax. The escape sequence for outer joins is:

```
{oj outer-join}
```

where *outer-join* is:

```
table-reference {LEFT | RIGHT | FULL} OUTER JOIN
```

```
{table-reference | outer-join} ON search-condition
```

where:

table-reference is a table name and *search-condition* is the join condition you want to use for the tables.

Example

```
SELECT Customers.CustID, Customers.Name, Orders.OrderID, Orders.Status
```

```
FROM {oj Customers LEFT OUTER JOIN
```

```
Orders ON Customers.CustID=Orders.CustID}
```

```
WHERE Orders.Status='OPEN'
```

The ADO.NET Data Provider supports the following outer join escape sequences as supported by Zen 9.x and higher:

- Left outer joins
- Right outer joins
- Full outer joins

D. Locking and Isolation Levels

Different database systems support various locking and isolation levels. The following topics cover locking and isolation levels and how their settings affect the data you retrieve:

- [Locking](#)
- [Isolation Levels](#)
- [Locking Modes and Levels](#)

Locking

Locking is a database operation that restricts a user from accessing a table or record. Locking is used in situations where more than one user might try to use the same table or record at the same time. By locking the table or record, the system ensures that only one user at a time can affect the data.

Locking is critical in multiuser databases, where different users can try to access or modify the same records concurrently. Although such concurrent database activity is desirable, it can create problems. Without locking, for example, if two users try to modify the same record at the same time, they might encounter problems ranging from retrieving bad data to deleting data that the other user needs. If, however, the first user to access a record can lock that record to temporarily prevent other users from modifying it, such problems can be avoided. Locking provides a way to manage concurrent database access while minimizing the various problems it can cause.

Isolation Levels

An isolation level represents a particular locking strategy employed in the database system to improve data consistency. The higher the isolation level, the more complex the locking strategy behind it. The isolation level provided by the database determines whether a transaction will encounter the following behaviors in data consistency:

Dirty reads	User 1 modifies a row. User 2 reads the same row before User 1 commits. User 1 performs a rollback. User 2 has read a row that has never really existed in the database. User 2 may base decisions on false data.
Non-repeatable reads	User 1 reads a row but does not commit. User 2 modifies or deletes the same row and then commits. User 1 rereads the row and finds it has changed (or has been deleted).
Phantom reads	User 1 uses a search condition to read a set of rows but does not commit. User 2 inserts one or more rows that satisfy this search condition, then commits. User 1 rereads the rows using the search condition and discovers rows that were not present before.

Isolation levels represent the database system's ability to prevent these behaviors. The American National Standards Institute (ANSI) defines four isolation levels:

- Read uncommitted (0)
- Read committed (1)
- Repeatable read (2)
- Serializable (3)

In ascending order (0–3), these isolation levels provide an increasing amount of data consistency to the transaction. At the lowest level, all three behaviors can occur. At the highest level, none can occur. The success of each level in preventing these behaviors is due to the locking strategies that they employ, which are as follows:

Read uncommitted (0)	Locks are obtained on modifications to the database and held until end of transaction (EOT). Reading from the database does not involve any locking.
Read committed (1)	Locks are acquired for reading and modifying the database. Locks are released after reading, but locks on modified objects are held until EOT.
Repeatable read (2)	Locks are obtained for reading and modifying the database. Locks on all modified objects are held until EOT. Locks obtained for reading data are held until EOT. Locks on non-modified access structures (such as indexes and hashing structures) are released after reading.

Serializable (3)	A lock is placed on the affected rows of the DataSet until EOT. All access structures that are modified, and those used by the query, are locked until EOT.
------------------	---

The following table shows what data consistency behaviors can occur at each isolation level.

Level	Dirty Read	Nonrepeatable Read	Phantom Read
0, Read uncommitted	Yes	Yes	Yes
1, Read committed	No	Yes	Yes
2, Repeatable read	No	No	Yes
3, Serializable	No	No	No

Although higher isolation levels provide better data consistency, this consistency can be costly in terms of the *concurrency* provided to individual users. Concurrency is the ability of multiple users to access and modify data simultaneously. As isolation levels increase, so does the chance that the locking strategy used will create problems in concurrency.

Put another way: The higher the isolation level, the more locking involved, and the more time users may spend waiting for data to be freed by another user. Because of this inverse relationship between isolation levels and concurrency, you must consider how people use the database before choosing an isolation level. You must weigh the trade-offs between data consistency and concurrency, and decide which is more important.

Locking Modes and Levels

Different database systems employ various locking modes, but they have two basic modes in common: shared and exclusive. Shared locks can be held on a single object by multiple users. If one user has a shared lock on a record, then a second user can also get a shared lock on that same record; however, the second user cannot get an exclusive lock on that record. Exclusive locks are exclusive to the user that obtains them. If one user has an exclusive lock on a record, then a second user cannot get either type of lock on the same record.

Performance and concurrency can also be affected by the locking level used in the database system. The locking level determines the size of an object that is locked in a database. For example, many database systems let you lock an entire table, as well as individual records. An intermediate level of locking, page-level locking, is also common. A page contains one or more records and is typically the amount of data read from the disk in a single disk access. The major disadvantage of page-level locking is that if one user locks a record, a second user may not be able to lock other records because they are stored on the same page as the locked record.

E. Designing .NET Applications for Performance Optimization

Developing performance-oriented .NET applications is not easy. Zen ADO.NET data providers do not throw exceptions to say that your code is running too slowly.

Retrieving Data

To retrieve data efficiently, return only the data that you need, and choose the most efficient method of doing so. The guidelines in this section will help you to optimize system performance when retrieving data with .NET applications.

Retrieving Long Data

Unless it is necessary, applications should not request long data because retrieving long data across a network is slow and resource-intensive.

Most users don't want to see long data. If the user wants to see these result items, then the application can query the database again, specifying only the long columns in the select list. This method allows the average user to retrieve the result set without having to pay a high performance penalty for network traffic.

Although the best method is to exclude long data from the select list, some applications do not formulate the select list before sending the query to the Zen ADO.NET data provider (that is, some applications use syntax such as `select * from <table name> ...`). If the select list contains long data, then some data providers must retrieve that data at fetch time even if the application does not bind the long data in the result set. When possible, try to implement a method that does not retrieve all columns of the table.

Sometimes long data must be retrieved. When this is the case, remember that most users do not want to see 100 KB, or more, of text on the screen.

Reducing the Size of Data Retrieved

To reduce network traffic and improve performance, you can reduce the size of any data being retrieved to some manageable limit by calling `set max rows` or `set max field size`, or some other database-specific command to limit row size or field size. Another method of reducing the size of the data being retrieved is to decrease the column size. If the data provider allows you to define the packet size, use the smallest packet size that will meet your needs.

In addition, be careful to return only the rows you need. If you return five columns when you only need two columns, performance is decreased, especially if the unnecessary rows include long data.

Using CommandBuilder Objects

It is tempting to use CommandBuilder objects because they generate SQL statements. However, this shortcut can have a negative effect on performance. Because of concurrency restrictions, the Command Builder does not generate efficient SQL statements. For example, the following SQL statement was created by the Command Builder:

```
CommandText: UPDATE TEST01.EMP SET EMPNO = ?, ENAME = ?, JOB = ?, MGR = ?, HIREDATE = ?, SAL = ?, COMM  
= ?, DEPT = ?
```

```
WHERE  
( (EMPNO = ?) AND ((ENAME IS NULL AND ? IS NULL)  
OR (ENAME = ?)) AND ((JOB IS NULL AND ? IS NULL)  
OR (JOB = ?)) AND ((MGR IS NULL AND ? IS NULL)  
OR (MGR = ?)) AND ((HIREDATE IS NULL AND ? IS NULL)  
OR (HIREDATE = ?)) AND ((SAL IS NULL AND ? IS NULL)  
OR (SAL = ?)) AND ((COMM IS NULL AND ? IS NULL)  
OR (COMM = ?)) AND ((DEPT IS NULL AND ? IS NULL)  
OR (DEPT = ?)) )
```

The end user can often write more efficient update and delete statements than those that the CommandBuilder generates.

Another drawback is also implicit in the design of the CommandBuilder object. The CommandBuilder object is always associated with a DataAdapter object and registers itself as a listener for RowUpdating and RowUpdated events that the DataAdapter object generates. This means that two events must be processed for every row that is updated.

Choosing the Right Data Type

Retrieving and sending certain data types can be expensive. When you design a schema, select the data type that can be processed most efficiently. For example, integer data is processed faster than floating-point data. Floating-point data is defined according to internal database-specific formats, usually in a compressed format. The data must be decompressed and converted into a different format so that it can be processed by the wire protocol.

Processing time is shortest for character strings, followed by integers, which usually require some conversion or byte ordering. Processing floating-point data and timestamps is at least twice as slow as integers.

Selecting .NET Objects and Methods

The guidelines in this section will help you to optimize system performance when selecting and using .NET objects and methods.

Using Parameter Markers as Arguments to Stored Procedures

When calling stored procedures, always use parameter markers for the argument markers instead of using literal arguments.

When you set the CommandText property in the Command object to the stored procedure name, do not physically code the literal arguments into the CommandText. For example, do not use literal arguments such as:

```
{call expense (3567, 'John', 987.32)}
```

Zen ADO.NET data providers can call stored procedures on the database server by executing the procedure as any other SQL query. Executing the stored procedure as a SQL query results in the database server parsing the statement, validating the argument types, and converting the arguments into the correct data types.

In the following example, the application programmer might assume that the only argument to getCustName is the integer 12345:

```
{call getCustName (12345)}
```

However, SQL is always sent to the database server as a character string. When the database server parses the SQL query and isolates the argument value, the result is still a string. The database server must then convert the string '12345' into the integer 12345. Using a parameter marker eliminates the need to convert the string and reduces the amount of processing by the server:

```
{call getCustName (?)}
```

Designing .NET Applications

The guidelines in this section will help you to optimize system performance when designing .NET applications.

Managing Connections

Connection management is important to application performance. Optimize your application by connecting once and using multiple statement objects, instead of performing multiple connections. Avoid connecting to a data source after establishing an initial connection.

You can improve performance significantly with connection pooling, especially for applications that connect over a network or through the World Wide Web. Connection pooling lets you reuse connections. Closing connections does not close the physical connection to the database. When an application requests a connection, an active connection is reused, thus avoiding the network I/O needed to create a new connection.

Pre-allocate connections. Decide what connection strings you will need to meet your needs. Remember that each unique connection string creates a new connection pool.

Once created, connection pools are not destroyed until the active process ends or the connection lifetime is exceeded. Maintenance of inactive or empty pools involves minimal system overhead.

Connection and statement handling should be addressed before implementation. Spending time and thoughtfully handling connection management improves application performance and maintainability.

Opening and Closing Connections

Open connections just before they are needed. Opening them earlier than necessary decreases the number of connections available to other users and can increase the demand for resources.

To keep resources available, explicitly Close the connection as soon as it is no longer needed. If you wait for the garbage collector to implicitly clean up connections that go out of scope, the connections will not be returned to the connection pool immediately, tying up resources that are not actually being used.

Close connections inside a finally block. Code in the finally block always runs, even if an exception occurs. This guarantees explicit closing of connections. For example:

```
try
{
    DBConn.Open();
    ... // Do some other interesting work
```

```
}  
catch (Exception ex)  
{  
    // Handle exceptions  
}  
finally  
{  
    // Close the connection  
    if (DBConn != null)  
        DBConn.Close();  
}
```

If you are using connection pooling, opening and closing connections is not an expensive operation. Using the Close() method of the data provider's Connection object adds or returns the connection to the connection pool. Remember, however, that closing a connection automatically closes all DataReader objects that are associated with the connection.

Using Statement Caching

A statement cache is a group of prepared statements or instances of Command objects that can be reused by an application. Using statement caching can improve application performance because the actions on the prepared statement are performed once even though the statement is reused multiple times over an application's lifetime.

A statement cache is owned by a physical connection. After being executed, a prepared statement is placed in the statement cache and remains there until the connection is closed.

Caching all of the prepared statements that an application uses might appear to offer increased performance. However, this approach may come at a cost of database memory if you implement statement caching with connection pooling. In this case, each pooled connection has its own statement cache that may contain all of the prepared statements that are used by the application. All of these pooled prepared statements are also maintained in the database's memory.

Using Commands Multiple Times

Choosing whether to use the Command.Prepare method can have a significant positive (or negative) effect on query execution performance. The Command.Prepare method tells the underlying data provider to optimize for multiple executions of statements that use parameter markers. Note that it is possible to Prepare any command regardless of which execution method is used (ExecuteReader, ExecuteNonQuery, or ExecuteScalar).

Consider the case where a Zen ADO.NET data provider implements Command.Prepare by creating a stored procedure on the server that contains the prepared statement. Creating stored procedures involves substantial overhead, but the statement can be executed multiple times. Although creating stored procedures is performance-expensive, execution of that statement is

minimized because the query is parsed and optimization paths are stored at create procedure time. Applications that execute the same statement multiples times can benefit greatly from calling `Command.Prepare` and then executing that command multiple times.

However, using `Command.Prepare` for a statement that is executed only once results in unnecessary overhead. Furthermore, applications that use `Command.Prepare` for large single execution query batches exhibit poor performance. Similarly, applications that either always use `Command.Prepare` or never use `Command.Prepare` do not perform as well as those that use a logical combination of prepared and unprepared statements.

Using Native Managed Providers

Bridges into unmanaged code, that is, code outside the .NET environment, adversely affect performance. Calling unmanaged code from managed code causes the data provider to be significantly slower than data providers that are completely managed code. Why take that kind of performance hit?

If you use a bridge, your code will be written for this bridge. Later, when a database-specific Zen ADO.NET data provider becomes available, the code must be rewritten; you will have to rewrite object names, schema information, error handling, and parameters. You'll save valuable time and resources by coding to managed data providers instead of coding to the bridges.

Updating Data

This section provides general guidelines to help you to optimize system performance when updating data in databases.

Using the Disconnected DataSet

Keep result sets small. The full result set must be retrieved from the server before the DataSet is populated. The full result set is stored in memory on the client.

Synchronizing Changes Back to the Data Source

You must build the logic into the PsqlDataAdapter for synchronizing the changes back to the data source using the primary key, as shown in the following example:

```
string updateSQL As String = "UPDATE emp SET sal = ?, job = ?" +  
    " = WHERE empno = ?";
```

F. Using an .edmx File

An .edmx file is an XML file that defines an Entity Data Model (EDM), describes the target database schema, and defines the mapping between the EDM and the database. An .edmx file also contains information that is used by the ADO.NET Entity Data Model Designer (Entity Designer) to render a model graphically.

The following code examples illustrate the necessary changes to the .edmx file in order to provide Extended Entity Framework functionality to the EDM layer.

The Entity Framework includes a set of methods similar to those of ADO.NET. These methods have been tailored to be useful for the new Entity Framework consumers – LINQ, EntitySQL, and ObjectServices.

The ADO.NET Entity Framework data provider models this functionality in the EDM by surfacing the `PsqlStatus` and `PsqlConnectionStatistics` entities, allowing you to model this functionality using standard tools in Visual Studio.

Code Examples

The following code fragment is an example of the SSDL model:

```
<?xml version="1.0" encoding="utf-8"?>
<Schema Namespace="EntityModel.Store" Alias="Self"
  xmlns="http://schemas.microsoft.com/ado/2006/04/edm/ssdl"
  Provider="Pervasive.Data.SqlClient"
  ProviderManifestToken="PSQL">
  <EntityContainer Name="SampleStoreContainer">
    <EntitySet Name="Connection_Statistics"
      EntityType="EntityModel.Store.Connection_Statistics" />
    <EntitySet Name="Status" EntityType="EntityModel.Store.Status" />
  </EntityContainer>
  <Function Name="RetrieveStatistics" Aggregate="false" BuiltIn="false" NiladicFunction="false"
    IsComposable="false" ParameterTypeSemantics="AllowImplicitConversion"
    StoreFunctionName=""Psql_Connection_RetrieveStatistics"" />
  <Function Name="EnableStatistics" Aggregate="false" BuiltIn="false" NiladicFunction="false"
    IsComposable="false" ParameterTypeSemantics="AllowImplicitConversion"
    StoreFunctionName=""Psql_Connection_EnableStatistics"" />
  <Function Name="DisableStatistics" Aggregate="false" BuiltIn="false" NiladicFunction="false"
    IsComposable="false" ParameterTypeSemantics="AllowImplicitConversion"
    StoreFunctionName=""Psql_Connection_DisableStatistics"" />
  <Function Name="ResetStatistics" Aggregate="false" BuiltIn="false" NiladicFunction="false"
    IsComposable="false" ParameterTypeSemantics="AllowImplicitConversion"
    StoreFunctionName=""Psql_Connection_ResetStatistics"" />
  </Function>
  <EntityType Name="Connection_Statistics">
    <Key>
      <PropertyRef Name="Id" />
    </Key>
    <Property Name="SocketReadTime" Type="double" Nullable="false" />
    <Property Name="MaxSocketReadTime" Type="double" Nullable="false" />
    <Property Name="SocketReads" Type="bigint" Nullable="false" />
    <Property Name="BytesReceived" Type="bigint" Nullable="false" />
    <Property Name="MaxBytesPerSocketRead" Type="bigint" Nullable="false" />
    <Property Name="SocketWriteTime" Type="double" Nullable="false" />
    <Property Name="MaxSocketWriteTime" Type="double" Nullable="false" />
    <Property Name="SocketWrites" Type="bigint" Nullable="false" />
    <Property Name="BytesSent" Type="bigint" Nullable="false" />
    <Property Name="MaxBytesPerSocketWrite" Type="bigint" Nullable="false" />
    <Property Name="TimeToDisposeOfUnreadRows" Type="double" Nullable="false" />
    <Property Name="SocketReadsToDisposeUnreadRows" Type="bigint" Nullable="false" />
    <Property Name="BytesRecvToDisposeUnreadRows" Type="bigint" Nullable="false" />
    <Property Name="IDUCount" Type="bigint" Nullable="false" />
    <Property Name="SelectCount" Type="bigint" Nullable="false" />
    <Property Name="StoredProcedureCount" Type="bigint" Nullable="false" />
    <Property Name="DDLCount" Type="bigint" Nullable="false" />
    <Property Name="PacketsReceived" Type="bigint" Nullable="false" />
    <Property Name="PacketsSent" Type="bigint" Nullable="false" />
    <Property Name="ServerRoundTrips" Type="bigint" Nullable="false" />
    <Property Name="SelectRowsRead" Type="bigint" Nullable="false" />
    <Property Name="StatementCacheHits" Type="bigint" Nullable="false" />
    <Property Name="StatementCacheMisses" Type="bigint" Nullable="false" />
    <Property Name="StatementCacheReplaces" Type="bigint" Nullable="false" />
    <Property Name="StatementCacheTopHit1" Type="bigint" Nullable="false" />
    <Property Name="StatementCacheTopHit2" Type="bigint" Nullable="false" />
    <Property Name="StatementCacheTopHit3" Type="bigint" Nullable="false" />
    <Property Name="PacketsReceivedPerSocketRead" Type="double" Nullable="false" />
    <Property Name="BytesReceivedPerSocketRead" Type="double" Nullable="false" />
    <Property Name="PacketsSentPerSocketWrite" Type="double" Nullable="false" />
    <Property Name="BytesSentPerSocketWrite" Type="double" Nullable="false" />
    <Property Name="PacketsSentPerRoundTrip" Type="double" Nullable="false" />
```

```

        <Property Name="PacketsReceivedPerRoundTrip" Type="double" Nullable="false" />
        <Property Name="BytesSentPerRoundTrip" Type="double" Nullable="false" />
        <Property Name="BytesReceivedPerRoundTrip" Type="double" Nullable="false" />
        <Property Name="Id" Type="integer" Nullable="false" />
    </EntityType>
    <EntityType Name="Status">
        <Key>
            <PropertyRef Name="Id" />
        </Key>
        <Property Name="ServerVersion" Type="varchar" MaxLength="127" Nullable="false" />
        <Property Name="Host" Type="varchar" MaxLength="127" Nullable="false" />
        <Property Name="Port" Type="integer" Nullable="false" />
        <Property Name="DatabaseName" Type="varchar" MaxLength="127" Nullable="false" />
        <Property Name="SessionId" Type="integer" Nullable="false" />
        <Property Name="StatisticsEnabled" Type="smallint_as_boolean" Nullable="false" />
        <Property Name="Id" Type="integer" Nullable="false" />
    </EntityType>
</Schema>

```

The following code fragment is an example of the MSL model:

```

<?xml version="1.0" encoding="utf-8"?>
<Mapping Space="C-S" xmlns="urn:schemas-microsoft-com:windows:storage:mapping:CS">
  <EntityContainerMapping
    StorageEntityContainer="SampleStoreContainer"
    CdmEntityContainer="SampleContainer">
    <EntitySetMapping Name="PsqlConnectionStatistics">
      <EntityTypeMapping TypeName="EntityModel.PsqlConnectionStatistics">
        <MappingFragment StoreEntitySet="Connection_Statistics">
          <ScalarProperty Name="SocketReadTime" ColumnName="SocketReadTime" />
          <ScalarProperty Name="MaxSocketReadTime" ColumnName="MaxSocketReadTime" />
          <ScalarProperty Name="SocketReads" ColumnName="SocketReads" />
          <ScalarProperty Name="BytesReceived" ColumnName="BytesReceived" />
          <ScalarProperty Name="MaxBytesPerSocketRead" ColumnName="MaxBytesPerSocketRead" />
          <ScalarProperty Name="SocketWriteTime" ColumnName="SocketWriteTime" />
          <ScalarProperty Name="MaxSocketWriteTime" ColumnName="MaxSocketWriteTime" />
          <ScalarProperty Name="SocketWrites" ColumnName="SocketWrites" />
          <ScalarProperty Name="BytesSent" ColumnName="BytesSent" />
          <ScalarProperty Name="MaxBytesPerSocketWrite" ColumnName="MaxBytesPerSocketWrite" />
          <ScalarProperty Name="TimeToDisposeOfUnreadRows"
ColumnName="TimeToDisposeOfUnreadRows" />
          <ScalarProperty Name="SocketReadsToDisposeUnreadRows"
ColumnName="SocketReadsToDisposeUnreadRows" />
          <ScalarProperty Name="BytesRecvToDisposeUnreadRows"
ColumnName="BytesRecvToDisposeUnreadRows" />
          <ScalarProperty Name="IDUCount" ColumnName="IDUCount" />
          <ScalarProperty Name="SelectCount" ColumnName="SelectCount" />
          <ScalarProperty Name="StoredProcedureCount" ColumnName="StoredProcedureCount" />
          <ScalarProperty Name="DDLCount" ColumnName="DDLCount" />
          <ScalarProperty Name="PacketsReceived" ColumnName="PacketsReceived" />
          <ScalarProperty Name="PacketsSent" ColumnName="PacketsSent" />
          <ScalarProperty Name="ServerRoundTrips" ColumnName="ServerRoundTrips" />
          <ScalarProperty Name="SelectRowsRead" ColumnName="SelectRowsRead" />
          <ScalarProperty Name="StatementCacheHits" ColumnName="StatementCacheHits" />
          <ScalarProperty Name="StatementCacheMisses" ColumnName="StatementCacheMisses" />
          <ScalarProperty Name="StatementCacheReplaces" ColumnName="StatementCacheReplaces" />
          <ScalarProperty Name="StatementCacheTopHit1" ColumnName="StatementCacheTopHit1" />
          <ScalarProperty Name="StatementCacheTopHit2" ColumnName="StatementCacheTopHit2" />
          <ScalarProperty Name="StatementCacheTopHit3" ColumnName="StatementCacheTopHit3" />
          <ScalarProperty Name="PacketsReceivedPerSocketRead"
ColumnName="PacketsReceivedPerSocketRead" />
          <ScalarProperty Name="BytesReceivedPerSocketRead"
ColumnName="BytesReceivedPerSocketRead" />
          <ScalarProperty Name="PacketsSentPerSocketWrite"

```

```

ColumnName="PacketsSentPerSocketWrite" />
    <ScalarProperty Name="BytesSentPerSocketWrite" ColumnName="BytesSentPerSocketWrite" />
    <ScalarProperty Name="PacketsSentPerRoundTrip" ColumnName="PacketsSentPerRoundTrip" />
    <ScalarProperty Name="PacketsReceivedPerRoundTrip"
ColumnName="PacketsReceivedPerRoundTrip" />
    <ScalarProperty Name="BytesSentPerRoundTrip" ColumnName="BytesSentPerRoundTrip" />
    <ScalarProperty Name="BytesReceivedPerRoundTrip"
ColumnName="BytesReceivedPerRoundTrip" />
    <ScalarProperty Name="Id" ColumnName="Id" />
    </MappingFragment>
</EntityTypeMapping>
</EntitySetMapping>
<EntitySetMapping Name="PsqlStatus">
    <EntityTypeMapping TypeName="EntityModel.PsqlStatus">
        <MappingFragment StoreEntitySet="Status">
            <ScalarProperty Name="ServerVersion" ColumnName="ServerVersion" />
            <ScalarProperty Name="Host" ColumnName="Host" />
            <ScalarProperty Name="Port" ColumnName="Port" />
            <ScalarProperty Name="DatabaseName" ColumnName="DatabaseName" />
            <ScalarProperty Name="SessionId" ColumnName="SessionId" />
            <ScalarProperty Name="StatisticsEnabled" ColumnName="StatisticsEnabled" />
            <ScalarProperty Name="Id" ColumnName="Id" />
        </MappingFragment>
    </EntityTypeMapping>
</EntitySetMapping>
<FunctionImportMapping FunctionImportName="RetrieveStatistics"
FunctionName="EntityModel.Store.RetrieveStatistics" />
<FunctionImportMapping FunctionImportName="EnableStatistics"
FunctionName="EntityModel.Store.EnableStatistics" />
<FunctionImportMapping FunctionImportName="DisableStatistics"
FunctionName="EntityModel.Store.DisableStatistics" />
<FunctionImportMapping FunctionImportName="ResetStatistics"
FunctionName="EntityModel.Store.ResetStatistics" />
<FunctionImportMapping FunctionImportName="Reauthenticate"
FunctionName="EntityModel.Store.Reauthenticate" />
</EntityContainerMapping>
</Mapping>

```

Breaking the model down further, a CSDL model is established at the conceptual layer – this is what is exposed to the EDM.

```

<?xml version="1.0" encoding="utf-8"?>
<Schema Namespace="EntityModel" Alias="Self" xmlns="http://schemas.microsoft.com/ado/2006/04/edm">
    <EntityContainer Name="SampleContainer">
        <EntitySet Name="PsqlConnectionStatistics" EntityType="EntityModel.PsqlConnectionStatistics" />
        <EntitySet Name="PsqlStatus" EntityType="EntityModel.PsqlStatus" />
        <FunctionImport Name="RetrieveStatistics" EntitySet="PsqlConnectionStatistics"
ReturnType="Collection(EntityModel.PsqlConnectionStatistics)" />
        <FunctionImport Name="EnableStatistics" EntitySet="PsqlStatus"
ReturnType="Collection(EntityModel.PsqlStatus)" />
        <FunctionImport Name="DisableStatistics" EntitySet="PsqlStatus"
ReturnType="Collection(EntityModel.PsqlStatus)" />
        <FunctionImport Name="ResetStatistics" EntitySet="PsqlStatus"
ReturnType="Collection(EntityModel.PsqlStatus)" />
        <FunctionImport Name="Reauthenticate" EntitySet="PsqlStatus"
ReturnType="Collection(EntityModel.PsqlStatus)" />
        <Parameter Name="CurrentUser" Type="String" />
        <Parameter Name="CurrentPassword" Type="String" />
        <Parameter Name="CurrentUserAffinityTimeout" Type="Int32" />
    </FunctionImport>
</EntityContainer>
<EntityType Name="PsqlConnectionStatistics">
    <Key>
        <PropertyRef Name="Id" />
    </Key>

```

```

</Key>
<Property Name="SocketReadTime" Type="Double" Nullable="false" />
<Property Name="MaxSocketReadTime" Type="Double" Nullable="false" />
<Property Name="SocketReads" Type="Int64" Nullable="false" />
<Property Name="BytesReceived" Type="Int64" Nullable="false" />
<Property Name="MaxBytesPerSocketRead" Type="Int64" Nullable="false" />
<Property Name="SocketWriteTime" Type="Double" Nullable="false" />
<Property Name="MaxSocketWriteTime" Type="Double" Nullable="false" />
<Property Name="SocketWrites" Type="Int64" Nullable="false" />
<Property Name="BytesSent" Type="Int64" Nullable="false" />
<Property Name="MaxBytesPerSocketWrite" Type="Int64" Nullable="false" />
<Property Name="TimeToDisposeOfUnreadRows" Type="Double" Nullable="false" />
<Property Name="SocketReadsToDisposeUnreadRows" Type="Int64" Nullable="false" />
<Property Name="BytesRecvToDisposeUnreadRows" Type="Int64" Nullable="false" />
<Property Name="IDUCount" Type="Int64" Nullable="false" />
<Property Name="SelectCount" Type="Int64" Nullable="false" />
<Property Name="StoredProcedureCount" Type="Int64" Nullable="false" />
<Property Name="DDLCount" Type="Int64" Nullable="false" />
<Property Name="PacketsReceived" Type="Int64" Nullable="false" />
<Property Name="PacketsSent" Type="Int64" Nullable="false" />
<Property Name="ServerRoundTrips" Type="Int64" Nullable="false" />
<Property Name="SelectRowsRead" Type="Int64" Nullable="false" />
<Property Name="StatementCacheHits" Type="Int64" Nullable="false" />
<Property Name="StatementCacheMisses" Type="Int64" Nullable="false" />
<Property Name="StatementCacheReplaces" Type="Int64" Nullable="false" />
<Property Name="StatementCacheTopHit1" Type="Int64" Nullable="false" />
<Property Name="StatementCacheTopHit2" Type="Int64" Nullable="false" />
<Property Name="StatementCacheTopHit3" Type="Int64" Nullable="false" />
<Property Name="PacketsReceivedPerSocketRead" Type="Double" Nullable="false" />
<Property Name="BytesReceivedPerSocketRead" Type="Double" Nullable="false" />
<Property Name="PacketsSentPerSocketWrite" Type="Double" Nullable="false" />
<Property Name="BytesSentPerSocketWrite" Type="Double" Nullable="false" />
<Property Name="PacketsSentPerRoundTrip" Type="Double" Nullable="false" />
<Property Name="PacketsReceivedPerRoundTrip" Type="Double" Nullable="false" />
<Property Name="BytesSentPerRoundTrip" Type="Double" Nullable="false" />
<Property Name="BytesReceivedPerRoundTrip" Type="Double" Nullable="false" />
<Property Name="Id" Type="Int32" Nullable="false" />
</EntityType>
<EntityType Name="PsqlStatus">
  <Key>
    <PropertyRef Name="Id" />
  </Key>
  <Property Name="ServerVersion" Type="String" Nullable="false" />
  <Property Name="Host" Type="String" Nullable="false" />
  <Property Name="Port" Type="Int32" Nullable="false" />
  <Property Name="DatabaseName" Type="String" Nullable="false" />
  <Property Name="SessionId" Type="Int32" Nullable="false" />
  <Property Name="StatisticsEnabled" Type="Boolean" Nullable="false" />
  <Property Name="Id" Type="Int32" Nullable="false" />
</EntityType>
</Schema>

```


G. Bulk Load Configuration Files

The following topics describe the configuration files used by Zen Bulk Load.

- [Sample Bulk Data Configuration File](#)
- [XML Schema Definition for a Bulk Data Configuration File](#)

See [Using Zen Bulk Load](#) for more information about this feature.

Sample Bulk Data Configuration File

The bulk format configuration file is produced when either a table or a DataReader is exported (unloaded) using the BulkCopy and BulkLoad operation.

```
<?xml version="1.0"?>
<!--
Sample DDL
-----

CREATE_STMT = CREATE TABLE GTABLE (CHARCOL char(10),VCHARCOL varchar2(10), \
DECIMALCOL number(15,5), NUMERICCOL decimal(15,5), SMALLCOL number(38), \
INTEGERCOL integer, REALCOL number, \
FLOATCOL float, DOUBLECOL number, LVCOL clob, \
BITCOL number(1),TINYINTCOL number(19), BIGINTCOL number(38), BINCOL raw(10), \
VARBINCOL raw(10), LVARBINCOL blob, DATECOL date, \
TIMECOL date, TSCOL date) -->

<table codepage="UTF-16" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="file:///c:/Documents and Settings/jbruce/My Documents/Current Specs/
BulkData.xsd">
  <row>
    <column codepage="UTF-16" datatype="CHAR" length="10" nullable="true">CHARCOL</column>
    <column codepage="UTF-16" datatype="LONGVARCHAR" length="10">VCHARCOL</column>
    <column codepage="UTF-16" datatype="DECIMAL" precision="15" scale="5">DECIMALCOL</column>
    <column codepage="UTF-16" datatype="DECIMAL" precision="15" scale="5">NUMERICCOL</column>
    <column codepage="UTF-16" datatype="DECIMAL" precision="38">SMALLCOL</column>
    <column codepage="UTF-16" datatype="INTEGER">INTEGERCOL</column>
    <column codepage="UTF-16" datatype="SINGLEPRECISION">REALCOL</column>

    <!--More definitions can follow here -->
  </row>
</table>
```

XML Schema Definition for a Bulk Data Configuration File

The bulk configuration XML schema governs bulk configuration files. The bulk configuration files in turn govern bulk load data files that are processed by Zen Bulk Load.

This schema provides a foundation for any third-party functionality to be built using this standard. Custom applications or tools that manage large amounts of data can employ this schema as a loosely coupled Zen Bulk Load across ODBC, JDBC, and ADO.NET APIs and across multiple platforms.

If you want to generate CSV data that can be consumed by Zen Bulk Load, you must supply an XML Schema for your XML configuration file.

Each bulk operation generates an XML configuration file in UTF-8 format that describes the bulk data file produced. If the bulk data file cannot be created or does not comply with the schema described in the XML configuration file, an exception is returned.

H. IANA Code Page Mappings

This following table maps the most widely used IBM code pages to IANA code page names.

IBM Number	IANA Code Page Name
37	IBM037
38	IBM038
290	IBM290
300	IBM300
301	IBM301
500	IBM500
813	ISO_8859-7:1987
819	ISO_8859-1:1987
857	IBM857
860	IBM860
861	IBM861
897	IBM897
932	IBM-942_P120-2000
939	IBM-939
943	Windows-932-2000 (for Windows clients)
943	IBM-943_P14A-2000 (for UNIX clients)
950	Big5
1200	UTF-16
1208	UTF-8
1251	Windows-1251
1252	Windows-1252
4396	IBM-930
5025	IBM5025
5035	IBM5035

IBM Number	IANA Code Page Name
5297	UTF-16
5304	UTF-8
13488	UTF-16BE

I. Glossary

.NET Architecture

Microsoft defines Microsoft .NET as a set of software technologies for connecting information, people, systems, and devices. To optimize software integration, the .NET Framework uses small, discrete building-block applications called Web services that connect to each other as well as to other, larger applications over the Internet.

The .NET Framework has two key parts:

- ASP.NET is an environment for building smart client applications (Windows Forms), and a loosely coupled data access subsystem (ADO.NET).
- The common language runtime (CLR) is the core runtime engine for executing applications in the .NET Framework. You can think of the CLR as a safe area – a "sandbox" – inside of which your .NET code runs. Code that runs in the CLR is called managed code.

ADO.NET

The data access component for the .NET Framework. ADO.NET is made of a set of classes that are used for connecting to a database, providing access to relational data, XML, and application data, and retrieving results.

ADO.NET Entity Framework

An object-relational mapping (ORM) framework for the .NET Framework. Developers can use it to create data access applications by programming against a conceptual application model instead of programming directly against a relational storage schema. This model allows developers to decrease the amount of code that must be written and maintained in data-centric applications

assembly

A compiled representation of one or more classes. Each assembly is self-contained, that is, the assembly includes the metadata about the assembly as a whole. Assemblies can be private or shared.

- Private assemblies, which are used by a limited number of applications, are placed in the application folder or one of its subfolders. For example, even if the client has two different

applications that call a private assembly named formulas, each client application loads the correct assembly.

- Shared assemblies, which are available to multiple client applications, are placed in the Global Assembly Cache (GAC). Each shared assembly is assigned a strong name to handle name and version conflicts.

bulk load

Rows from the database client are sent to the database server in a continuous stream instead of in numerous smaller packets of data. Performance improves because far fewer network round trips are required.

client load balancing

A mechanism that distributes new connections in a computing environment so that no one server is overwhelmed with connection requests.

code access security (CAS)

A mechanism provided by the common language runtime through which managed code is granted permissions by security policy; permissions are enforced, limiting the operations that the code will be allowed to perform.

common language runtime (CLR)

The common language runtime (CLR) is the core runtime engine in the Microsoft .NET Framework. The CLR supplies services such as cross-language integration, code access security, object lifetime management, and debugging support. Applications that run in the CLR are sometimes said to be running "in the sandbox."

connection failover

A mechanism that allows an application to connect to an alternate, or backup, database server if the primary database server is unavailable, for example, because of a hardware failure or traffic overload.

connection pooling

The process by which connections can be reused rather than creating a new one every time the data provider needs to establish a connection to the underlying database.

connection retry

Connection retry defines the number of times the data provider attempts to connect to the primary and, if configured, alternate database servers after the initial unsuccessful connection attempt. Connection retry can be an important strategy for system recovery.

Data Access Application Block (DAAB)

A pre-defined code block that provides access to the most often used ADO.NET data access features. Applications can use the application block to pass data through application layers, and submit changed data back to the database.

destination table

In a Zen Bulk Load operation, the table on the database server into which the data is copied.

entity

An entity is an instance of an EntityType. It has a unique identity, independent existence, and forms the operational unit of consistency. An EntityType defines the principal data objects about which information has to be managed such as person, places, things or activities relevant to the application.

global assembly cache (GAC)

The part of the assembly cache that stores assemblies specifically installed to be shared by many applications on the computer. Applications deployed in the global assembly cache must have a strong name to handle name and version conflicts.

isolation level

A particular locking strategy employed in the database system to improve data consistency. The higher the isolation level number, the more complex the locking strategy behind it. The isolation level provided by the database determines how a transaction handles data consistency.

The American National Standards Institute (ANSI) defines four isolation levels:

- Read uncommitted (0)
- Read committed (1)
- Repeatable read (2)
- Serializable (3)

load balancing

See [client load balancing](#).

locking level

A database operation that restricts a user from accessing a table or record. Locking is used in situations more than one user might try to use the same table or the same time. By locking the table or record, the system that only one user at a time can affect the data.

Logging Application Block (LAB)

A component of the Microsoft Enterprise Libraries that simplifies the implementation of common logging functions. Developers can use the Logging Block to write information to a variety of locations, such as the event log, an e-mail message, or a database.

managed code

Code executed and managed by the .NET Framework, specifically by the CLR. Managed code must supply the information necessary for the CLR to provide services such as memory management and code access security.

namespace

A logical naming scheme for grouping related types. The .NET Framework uses a hierarchical naming scheme for grouping types into logical categories of related functionality, such as the ASP.NET technology or remoting functionality. Design tools can use namespaces to make it easier for developers to browse and reference types in their code. A single assembly can contain types whose hierarchical names have different namespace roots, and a logical namespace root can span multiple assemblies. In the .NET Framework, a namespace is a logical design-time naming convenience, whereas an assembly establishes the name scope for types at run time.

Performance Monitor

A tool in the Windows SDK that identifies areas in which performance problems exist. You can use Performance Monitor to identify the number and frequency of CLR exceptions in your applications. The Performance Monitor (PerfMon) and VS Performance Monitor (VSPerfMon) utilities also allow you to record application parameters and review the results as a report or graph.

stream

An abstraction of a sequence of binary or text data. The Stream class and its derived classes provide a generic view of these different types of input and output.

schema collection

Closely related schemas that can be handled more efficiently when grouped together. Database schema elements such as tables and columns are exposed through schema collections.

strong name

A name that consists of an assembly's text name, version number, and culture information (if provided), with a public key and a digital signature generated over the assembly. Assemblies with the same strong name should be identical.

unmanaged code

Code that is executed directly by the operating system, outside of the CLR. Unmanaged code includes all code written before the .NET Framework was introduced. Because it is outside the .NET environment, unmanaged code cannot make use of any .NET managed facilities.
