

Interview Questions .NET Windows Forms

1. **Write a simple Windows Forms MessageBox statement.**
2. `System.Windows.Forms.MessageBox.Show ("Hello, Windows Forms");`

Can you write a class without specifying namespace? Which namespace does it belong to by default??

Yes, you can, then the class belongs to global namespace which has no name. For commercial products, naturally, you wouldn't want global namespace.

You are designing a GUI application with a window and several widgets on it. The user then resizes the app window and sees a lot of grey space, while the widgets stay in place. What's the problem? One should use anchoring for correct resizing. Otherwise the default property of a widget on a form is top-left, so it stays at the same location when resized.

How can you save the desired properties of Windows Forms application? .config files in .NET are supported through the API to allow storing and retrieving information. They are nothing more than simple XML files, sort of like what .ini files were before for Win32 apps.

So how do you retrieve the customized properties of a .NET application from XML .config file? Initialize an instance of AppSettingsReader class. Call the GetValue method of AppSettingsReader class, passing in the name of the property and the type expected. Assign the result to the appropriate variable.

Can you automate this process? In Visual Studio yes, use Dynamic Properties for automatic .config creation, storage and retrieval.

My progress bar freezes up and dialog window shows blank, when an intensive background process takes over. Yes, you should've multi-threaded your GUI, with taskbar and main form being one thread, and the background process being the other.

What's the safest way to deploy a Windows Forms app? Web deployment: the user always downloads the latest version of the code; the program runs within security sandbox, properly written app will not require additional security privileges.

Why is it not a good idea to insert code into InitializeComponent method when working with Visual Studio? The designer will likely throw it away; most of the code inside InitializeComponent is auto-generated.

What's the difference between WindowsDefaultLocation and WindowsDefaultBounds? WindowsDefaultLocation tells the form to start up at a location selected by OS, but with internally specified size. WindowsDefaultBounds delegates both size and starting position choices to the OS.

What's the difference between Move and LocationChanged?

Resize and SizeChanged? Both methods do the same, Move and Resize are the names adopted from VB to ease migration to C#.

How would you create a non-rectangular window, let's say an ellipse?

Create a rectangular form, set the TransparencyKey property to the same value as BackColor, which will effectively make the background of the form transparent. Then set the FormBorderStyle to FormBorderStyle.None, which will remove the contour and contents of the form.

How do you create a separator in the Menu Designer? A hyphen '-' would do it. Also, an ampersand '&\' would underline the next letter.

How's anchoring different from docking? Anchoring treats the component as having the absolute size and adjusts its location relative to the parent form. Docking treats the component location as absolute and disregards the component size. So if a status bar must always be at the bottom no matter what, use docking. If a button should be on the top right, but change its position with the form being resized, use anchoring.

Interview Questions

C#

What's the implicit name of the parameter that gets passed into the class' set method?

Value, and its datatype depends on whatever variable we're changing.

How do you inherit from a class in C#? Place a colon and then the name of the base class. Notice that it's double colon in C++.

Does C# support multiple inheritance? No, use interfaces instead.

When you inherit a protected class-level variable, who is it available to?

Classes in the same namespace.

Are private class-level variables inherited? Yes, but they are not accessible, so looking at it you can honestly say that they are not inherited. But they are.

Describe the accessibility modifier protected internal. It's available to derived classes and classes within the same Assembly (and naturally from the base class it's declared in).

C# provides a default constructor for me. I write a constructor that takes a string as a parameter, but want to keep the no parameter one. How many constructors should I write?

Two. Once you write at least one constructor, C# cancels the freebie constructor, and now you have to write one yourself, even if there's no implementation in it.

What's the top .NET class that everything is derived from?

System.Object.

How's method overriding different from overloading? When overriding, you change the method behavior for a derived class.

Overloading simply involves having a method with the same name within the class.

What does the keyword virtual mean in the method definition?

The method can be over-ridden.

Can you declare the override method static while the original method is non-static? No, you can't, the signature of the virtual method must remain the same, only the keyword virtual is changed to keyword override.

Can you override private virtual methods? No, moreover, you cannot access private methods in inherited classes, have to be protected in the base class to allow any sort of access.

Can you prevent your class from being inherited and becoming a base class for some other classes? Yes, that's what keyword sealed in the class definition is for. The developer trying to derive from your class will get a message: cannot inherit from Sealed class WhateverBaseClassName. It's the same concept as final class in Java.

Can you allow class to be inherited, but prevent the method from being over-ridden? Yes, just leave the class public and make the method sealed.

What's an abstract class? A class that cannot be instantiated. A concept in C++ known as pure virtual method. A class that must be inherited and have the methods over-ridden. Essentially, it's a blueprint for a class without any implementation.

When do you absolutely have to declare a class as abstract (as opposed to free-willed educated choice or decision based on UML diagram)? When at least one of the methods in the class is abstract. When the class itself is inherited from an abstract class, but not all base abstract methods have been over-ridden.

What's an interface class? It's an abstract class with public abstract methods all of which must be implemented in the inherited classes.

Why can't you specify the accessibility modifier for methods inside the interface? They all must be public. Therefore, to prevent you from getting the false impression that you have any freedom of choice, you are not allowed to specify any accessibility, it's public by default.

Can you inherit multiple interfaces?

Yes, why not.

And if they have conflicting method names? It's up to you to implement the method inside your own class, so implementation is left entirely up to you. This might cause a problem on a higher-level scale if similarly named methods from different interfaces expect different data, but as far as compiler cares you're okay.

What's the difference between an interface and abstract class? In the interface all methods must be abstract; in the abstract class some methods can be concrete. In the interface no accessibility modifiers are allowed, which is ok in abstract classes.

How can you overload a method? Different parameter data types, different number of parameters, different order of parameters.

If a base class has a bunch of overloaded constructors, and an inherited class has another bunch of overloaded constructors, can you enforce a call from an inherited constructor to an arbitrary base constructor? Yes, just place a colon, and then keyword base (parameter list to invoke the appropriate constructor) in the overloaded constructor definition inside the inherited class.

What's the difference between System.String and System.StringBuilder classes? System.String is immutable; System.StringBuilder was designed with the purpose of having a mutable string where a variety of operations can be performed.

What's the advantage of using System.Text.StringBuilder over System.String? StringBuilder is more efficient in the cases, where a lot of manipulation is done to the text. Strings are immutable, so each time it's being operated on, a new instance is created.

Can you store multiple data types in System.Array? No.

What's the difference between the System.Array.CopyTo() and System.Array.Clone()? The first one performs a deep copy of the array, the second one is shallow.

How can you sort the elements of the array in descending order? By calling Sort() and then Reverse() methods.

What's the .NET datatype that allows the retrieval of data by a unique key? Hashtable.

What's class SortedList underneath? A sorted Hashtable.

Will finally block get executed if the exception had not occurred? Yes.

What's the C# equivalent of C++ catch (...), which was a catch-all statement for any possible exception? A catch block that catches the exception of type System.Exception. You can also omit the parameter data type in this case and just write catch {}.

Can multiple catch blocks be executed? No, once the proper catch code fires off, the control is transferred to the finally block (if there are any), and then whatever follows the finally block.

Why is it a bad idea to throw your own exceptions? Well, if at that point you know that an error has occurred, then why not write the proper code to handle that error instead of passing a new Exception object to the catch block? Throwing your own exceptions signifies some design flaws in the project.

What's a delegate? A delegate object encapsulates a reference to a method. In C++ they were referred to as function pointers.

What's a multicast delegate? It's a delegate that points to and eventually fires off several methods.

How's the DLL Hell problem solved in .NET? Assembly versioning allows the application to specify not only the library it needs to run (which was available under Win32), but also the version of the assembly.

What are the ways to deploy an assembly? An MSI installer, a CAB archive, and XCOPY command.

What's a satellite assembly? When you write a multilingual or multi-cultural application in .NET, and want to distribute the core application

separately from the localized modules, the localized assemblies that modify the core application are called satellite assemblies.

What namespaces are necessary to create a localized application? System.Globalization, System.Resources.

What's the difference between // comments, /* */ comments and /// comments? Single-line, multi-line and XML documentation comments.

How do you generate documentation from the C# file commented properly with a command-line compiler? Compile it with a /doc switch.

What's the difference between <c> and <code> XML documentation tag? Single line code example and multiple-line code example.

Is XML case-sensitive? Yes, so <Student> and <student> are different elements.

What debugging tools come with the .NET SDK? CorDBG – command-line debugger, and DbgCLR – graphic debugger. Visual Studio .NET uses the DbgCLR. To use CorDbg, you must compile the original C# file using the /debug switch.

What does the This window show in the debugger? It points to the object that's pointed to by this reference. Object's instance data is shown.

What does assert() do? In debug compilation, assert takes in a Boolean condition as a parameter, and shows the error dialog if the condition is false. The program proceeds without any interruption if the condition is true.

What's the difference between the Debug class and Trace class? Documentation looks the same. Use Debug class for debug builds, use Trace class for both debug and release builds.

Why are there five tracing levels in

System.Diagnostics.TraceSwitcher? The tracing dumps can be quite verbose and for some applications that are constantly running you run the risk of overloading the machine and the hard drive there. Five levels range from None to Verbose, allowing to fine-tune the tracing activities.

Where is the output of TextWriterTraceListener redirected? To the Console or a text file depending on the parameter passed to the constructor.

How do you debug an ASP.NET Web application? Attach the aspnet_wp.exe process to the DbgClr debugger.

What are three test cases you should go through in unit testing? Positive test cases (correct data, correct output), negative test cases (broken or missing data, proper handling), exception test cases (exceptions are thrown and caught properly).

Can you change the value of a variable while debugging a C# application? Yes, if you are debugging via Visual Studio.NET, just go to Immediate window.

Explain the three services model (three-tier application).

Presentation (UI), business (logic and underlying code) and data (from storage or other sources).

What are advantages and disadvantages of Microsoft-provided data provider classes in ADO.NET? SQLServer.NET data provider is high-speed and robust, but requires SQL Server license purchased from Microsoft. OLE-DB.NET is universal for accessing other sources, like Oracle, DB2, Microsoft Access and Informix, but it's a .NET layer on top of OLE layer, so not the fastest thing in the world. ODBC.NET is a deprecated layer provided for backward compatibility to ODBC engines.

What's the role of the DataReader class in ADO.NET connections? It returns a read-only dataset from the data source when the command is executed.

What is the wildcard character in SQL? Let's say you want to query database with LIKE for all employees whose name starts with La. The wildcard character is %, the proper query with LIKE would involve 'La%'.

Explain ACID rule of thumb for transactions. Transaction must be Atomic (it is one unit of work and does not dependent on previous and following transactions), Consistent (data is either committed or roll back, no "in-between" case where something has been updated and something hasn't), Isolated (no transaction sees the intermediate results of the current transaction), Durable (the values persist if the data had been committed even if the system crashes right after).

What connections does Microsoft SQL Server support? Windows Authentication (via Active Directory) and SQL Server authentication (via Microsoft SQL Server username and passwords).

Which one is trusted and which one is untrusted? Windows Authentication is trusted because the username and password are checked with the Active Directory, the SQL Server authentication is untrusted, since SQL Server is the only verifier participating in the transaction.

Why would you use untrusted verification? Web Services might use it, as well as non-Windows applications.

What does the parameter Initial Catalog define inside Connection String? The database name to connect to.

What's the data provider name to connect to Access database? Microsoft.Access.

What does Dispose method do with the connection object? Deletes it from the memory.

What is a pre-requisite for connection pooling? Multiple processes must agree that they will share the same connection, where every parameter is the same, including the security settings.

Interview Questions .NET Remoting

What's a Windows process? It's an application that's running and had been allocated memory.

What's typical about a Windows process in regards to memory allocation? Each process is allocated its own block of available RAM space, no process can access another process' code or data. If the process crashes, it dies alone without taking the entire OS or a bunch of other applications down.

Why do you call it a process? What's different between process and application in .NET, not common computer usage, terminology? A process is an instance of a running application. An application is an executable on the hard drive or network. There can be numerous processes launched of the same application (5 copies of Word running), but 1 process can run just 1 application.

What distributed process frameworks outside .NET do you know? Distributed Computing Environment/Remote Procedure Calls (DEC/RPC), Microsoft Distributed Component Object Model (DCOM), Common Object Request Broker Architecture (CORBA), and Java Remote Method Invocation (RMI).

What are possible implementations of distributed applications in .NET? .NET Remoting and ASP.NET Web Services. If we talk about the Framework Class Library, noteworthy classes are in System.Runtime.Remoting and System.Web.Services.

When would you use .NET Remoting and when Web services? Use remoting for more efficient exchange of information when you control both ends of the application. Use Web services for open-protocol-based information exchange when you are just a client or a server with the other end belonging to someone else.

What's a proxy of the server object in .NET Remoting? It's a fake copy of the server object that resides on the client side and behaves as if it was the server. It handles the communication between real server object and the client object. This process is also known as **marshaling**.

What are remotable objects in .NET Remoting? Remotable objects are the objects that can be marshaled across the application domains. You can marshal by value, where a deep copy of the object is created and then passed to the receiver. You can also marshal by reference, where just a reference to an existing object is passed.

What are channels in .NET Remoting? Channels represent the objects that transfer the other serialized objects from one application domain to another and from one computer to another, as well as one process to another on the same box. A channel must exist before an object can be transferred.

What security measures exist for .NET Remoting in System.Runtime.Remoting? None. Security should be taken care of at the application level. Cryptography and other security techniques can be applied at application or server level.

What is a formatter? A formatter is an object that is responsible for encoding and serializing data into messages on one end, and deserializing and decoding messages into data on the other end.

Choosing between HTTP and TCP for protocols and Binary and SOAP for formatters, what are the trade-offs? Binary over TCP is the most efficient, SOAP over HTTP is the most interoperable.

What's SingleCall activation mode used for? If the server object is instantiated for responding to just one single request, the request should be made in SingleCall mode.

What's Singleton activation mode? A single object is instantiated regardless of the number of clients accessing it. Lifetime of this object is determined by lifetime lease.

How do you define the lease of the object? By implementing ILease interface when writing the class code.

Can you configure a .NET Remoting object via XML file? Yes, via machine.config and application level .config file (or web.config in ASP.NET). Application-level XML settings take precedence over machine.config.

How can you automatically generate interface for the remotable object in .NET with Microsoft tools? Use the Soapsuds tool.

Interview Questions ASP.NET

Describe the role of *inetinfo.exe*, *aspnet_isapi.dll*

and *aspnet_wp.exe* in the page loading process. *inetinfo.exe* is the Microsoft IIS server running, handling ASP.NET requests among other things. When an ASP.NET request is received (usually a file with .aspx extension), the ISAPI filter *aspnet_isapi.dll* takes care of it by passing the request to the actual worker process *aspnet_wp.exe*.

What's the difference between `Response.Write()`

and `Response.Output.Write()`? The latter one allows you to write formatted output.

What methods are fired during the page load? `Init()` - when the page is instantiated, `Load()` - when the page is loaded into server memory, `PreRender()` - the brief moment before the page is displayed to the user as HTML, `Unload()` - when page finishes loading.

Where does the Web page belong in the .NET Framework class hierarchy? `System.Web.UI.Page`

Where do you store the information about the user's locale?

`System.Web.UI.Page.Culture`

What's the difference between `CodeBehind="MyCode.aspx.cs"` and `Src="MyCode.aspx.cs"`? `CodeBehind` is relevant to Visual Studio.NET only.

What's a bubbled event? When you have a complex control, like `DataGrid`, writing an event processing routine for each object (cell, button, row, etc.) is quite tedious. The controls can bubble up their

eventhandlers, allowing the main DataGrid event handler to take care of its constituents.

Suppose you want a certain ASP.NET function executed on MouseOver over a certain button. Where do you add an event handler?

It's the Attributes property, the Add function inside that property. So btnSubmit.Attributes.Add("onMouseOver","someClientCode();")

What data type does the RangeValidator control support?

Integer,String and Date.

Explain the differences between Server-side and Client-side code?

Server-side code runs on the server. Client-side code runs in the clients' browser.

What type of code (server or client) is found in a Code-Behind class?

Server-side code.

Should validation (did the user enter a real date) occur server-side or client-side? Why? Client-side. This reduces an additional request to the server to validate the users input.

What does the "EnableViewState" property do? Why would I want it on or off?

It enables the viewstate on the page. It allows the page to save the users input on a form.

What is the difference between Server.Transfer and Response.Redirect? Why would I choose one over the other?

Server.Transfer is used to post a form to another page.

Response.Redirect is used to redirect the user to another page or site.

15.Can you explain the difference between an ADO.NET Dataset and an ADO Recordset?

- A DataSet can represent an entire relational database in memory, complete with tables, relations, and views.
- A DataSet is designed to work without any continuing connection to the original data source.
- Data in a DataSet is bulk-loaded, rather than being loaded on demand.
- There's no concept of cursor types in a DataSet.
- DataSets have no current record pointer You can use For Each loops to move through the data.
- You can store many edits in a DataSet, and write them to the original data source in a single operation.
- Though the DataSet is universal, other objects in ADO.NET come in different versions for different data sources.

Can you give an example of what might be best suited to place in the Application_Start and Session_Start subroutines? This is where you can set the specific variables for the Application and Session objects.

If I'm developing an application that must accommodate multiple security levels though secure login and my ASP.NET web application is spanned across three web-servers (using round-robin load balancing) what would be the best approach

to maintain login-in state for the users? Maintain the login state security through a database.

Can you explain what inheritance is and an example of when you might use it? When you want to inherit (use the functionality of) another class. Base Class Employee. A Manager class could be derived from the Employee base class.

Whats an assembly? Assemblies are the building blocks of the .NET framework. [Overview of assemblies from MSDN](#)

Describe the difference between inline and code behind. Inline code written along side the html in a page. Code-behind is code written in a separate file and referenced by the .aspx page.

Explain what a diffgram is, and a good use for one? The DiffGram is one of the two XML formats that you can use to render DataSet object contents to XML. For reading database data to an XML file to be sent to a Web Service.

Whats MSIL, and why should my developers need an appreciation of it if at all? MSIL is the Microsoft Intermediate Language. All .NET compatible languages will get converted to MSIL.

Which method do you invoke on the DataAdapter control to load your generated dataset with data? The .Fill() method

Can you edit data in the Repeater control? No, it just reads the information from its data source

Which template must you provide, in order to display data in a Repeater control? ItemTemplate

How can you provide an alternating color scheme in a Repeater control? Use the AlternatingItemTemplate

What property must you set, and what method must you call in your code, in order to bind the data from some data source to the Repeater control? You must set the DataSource property and call the DataBind method.

What base class do all Web Forms inherit from? The Page class.

Name two properties common in every validation control?

ControlToValidate property and Text property.

What tags do you need to add within the asp:datagrid tags to bind columns manually? Set AutoGenerateColumns Property to false on the datagrid tag

What tag do you use to add a hyperlink column to the DataGrid?
<asp:HyperLinkColumn>

What is the transport protocol you use to call a Web service?

SOAP is the preferred protocol.

True or False: A Web service can only be written in .NET? False

What does WSDL stand for? (Web Services Description Language)

Where on the Internet would you look for Web services?

(<http://www.uddi.org>)

Which property on a Combo Box do you set with a column name, prior to setting the DataSource, to display data in the combo box? DataTextField property

Which control would you use if you needed to make sure the values in two different controls matched? CompareValidator Control

True or False: To test a Web service you must create a windows application or Web application to consume this service? False, the webservice comes with a test page and it provides HTTP-GET method to test.

C#, .NET, XML, IIS - Interview Questions

- Framework
- OOPS
- C# Language features
- Access specifiers
- Constructor
- ADO.NET
- Asp.Net
- WebService & Remoting
- COM
- XML
- IIS
- Controls
- Programming

What is .NET Framework?

The .NET Framework has two main components: the common language runtime and the .NET Framework class library.

You can think of the runtime as an agent that manages code at execution time, providing core services such as memory management, thread management, and remoting, while also enforcing strict type safety and other forms of code accuracy that ensure security and robustness.

The class library, is a comprehensive, object-oriented collection of reusable types that you can use to develop applications ranging from traditional command-line or graphical user interface (GUI) applications to applications based on the latest innovations provided by ASP.NET, such as Web Forms and XML Web services.

What is CLR, CTS, CLS?

The .NET Framework provides a runtime environment called the Common Language Runtime or CLR (similar to the Java Virtual Machine or JVM in Java), which handles the execution of code and provides useful services for the implementation of the program. CLR takes care of code management at program execution and provides various beneficial services such as memory management, thread management, security management, code verification, compilation, and other system services. The managed code that targets CLR benefits from useful features such as cross-language integration, cross-language exception handling, versioning, enhanced security, deployment support, and debugging.

Common Type System (CTS) describes how types are declared, used and managed in the runtime and facilitates cross-language integration, type safety, and high performance code execution.

The CLS is simply a specification that defines the rules to support language integration in such a way that programs written in any language, yet can interoperate with one another, taking full advantage of inheritance, polymorphism, exceptions, and other features. These rules and the specification are documented in the ECMA proposed standard document, "Partition I Architecture", <http://msdn.microsoft.com/net/ecma/>

What are the new features of Framework 1.1 ?

Native Support for Developing Mobile Web Applications

Enable Execution of Windows Forms Assemblies Originating from the Internet Assemblies originating from the Internet zone—for example, Microsoft Windows® Forms controls embedded in an Internet-based Web page or Windows Forms assemblies hosted on an Internet Web server and loaded either through the Web browser or programmatically using the System.Reflection.Assembly.LoadFrom() method—now receive sufficient permission to execute in a semi-trusted manner. Default security policy has been changed so that assemblies assigned by the common language runtime (CLR) to the Internet zone code group now receive the constrained permissions associated with the Internet permission set. In the .NET Framework 1.0 Service Pack 1 and Service Pack 2, such applications received the permissions associated with the Nothing permission set and could not execute.

Enable Code Access Security for ASP.NET Applications

Systems administrators can now use code access security to further lock down the permissions granted to ASP.NET Web applications and Web services.

Although the operating system account under which an application runs imposes security restrictions on the application, the code access security system of the CLR can enforce additional restrictions on selected application resources based on policies specified by systems administrators. You can use this feature in a shared server environment (such as an Internet service provider (ISP) hosting multiple Web applications on one server) to isolate separate applications from one another, as well as with stand-alone servers where you want applications to run with the minimum necessary privileges.

Native Support for Communicating with ODBC and Oracle Databases

5. Unified Programming Model for Smart Client Application Development

The Microsoft .NET Compact Framework brings the CLR, Windows Forms controls, and other .NET Framework features to small devices. The .NET Compact Framework supports a large subset of the .NET Framework class library optimized for small devices.

6. Support for IPv6

The .NET Framework 1.1 supports the emerging update to the Internet Protocol, commonly referred to as IP version 6, or simply IPv6. This protocol is designed to significantly increase the address space used to identify communication endpoints in the Internet to accommodate its ongoing growth.

<http://msdn.microsoft.com/netframework/technologyinfo/Overview/whatsnew.aspx>

Is .NET a runtime service or a development platform?

Ans: It's both and actually a lot more. Microsoft .NET includes a new way of delivering software and services to businesses and consumers. A part of Microsoft.NET is the .NET Frameworks. The .NET frameworks SDK consists of two parts: the .NET common language runtime and the .NET class library. In addition, the SDK also includes command-line compilers for C#, C++, JScript, and VB. You use these compilers to build applications and components. These components require the runtime to execute so this is a development platform.

What is MSIL, IL?

When compiling to managed code, the compiler translates your source code into Microsoft intermediate language (MSIL), which is a CPU-independent set of instructions that can be efficiently converted to native code. MSIL includes instructions for loading, storing, initializing, and calling methods on objects, as well as instructions for arithmetic and logical operations, control flow, direct memory access, exception handling, and other operations. Microsoft intermediate language (MSIL) is a language used as the output of a number of compilers and as the input to a just-in-time (JIT) compiler. The common language runtime includes a JIT compiler for converting MSIL to native code.

Can I write IL programs directly?

Yes. Peter Drayton posted this simple example to the DOTNET mailing list:

```
.assembly MyAssembly {}  
.class MyApp {
```

```

.method static void Main() {
    .entrypoint
    ldstr    "Hello, IL!"
    call    void System.Console::WriteLine(class System.Object)
    ret
}
}

```

Just put this into a file called hello.il, and then run `ilasm hello.il`. An exe assembly will be generated.

Can I do things in IL that I can't do in C#?

Yes. A couple of simple examples are that you can throw exceptions that are not derived from `System.Exception`, and you can have non-zero-based arrays.

What is JIT (just in time)? how it works?

Before Microsoft intermediate language (MSIL) can be executed, it must be converted by a .NET Framework just-in-time (JIT) compiler to native code, which is CPU-specific code that runs on the same computer architecture as the JIT compiler.

Rather than using time and memory to convert all the MSIL in a portable executable (PE) file to native code, it converts the MSIL as it is needed during execution and stores the resulting native code so that it is accessible for subsequent calls.

The runtime supplies another mode of compilation called install-time code generation. The install-time code generation mode converts MSIL to native code just as the regular JIT compiler does, but it converts larger units of code at a time, storing the resulting native code for use when the assembly is subsequently loaded and executed.

As part of compiling MSIL to native code, code must pass a verification process unless an administrator has established a security policy that allows code to bypass verification. Verification examines MSIL and metadata to find out whether the code can be determined to be type safe, which means that it is known to access only the memory locations it is authorized to access.

What is strong name?

A name that consists of an assembly's identity—its simple text name, version number, and culture information (if provided)—strengthened by a public key and a digital signature generated over the assembly.

What is portable executable (PE)?

The file format defining the structure that all executable files (EXE) and Dynamic Link Libraries (DLL) must use to allow them to be loaded and executed by Windows. PE is derived from the Microsoft Common Object File Format (COFF). The EXE and DLL files created using the .NET Framework obey the PE/COFF formats and also add additional header and data sections to the files that are only used by the CLR. The specification for the PE/COFF file formats is available at

<http://www.microsoft.com/whdc/hwdev/hardware/pecoffdown.mspx>

What is Event - Delegate? clear syntax for writing a event delegate

```

12. {
13.     public static void Main()
14.     {
15.         TakesADelegate(new MyDelegate(DelegateFunction));
16.     }
17.     public static void TakesADelegate(MyDelegate SomeFunction)
18.     {
19.         SomeFunction(21);
20.     }
21.     public static void DelegateFunction(int i)
22.     {
23.         System.Console.WriteLine("Called by delegate with number: {0}.", i);
24.     }
}

```

What is Code Access Security (CAS)?

CAS is the part of the .NET security model that determines whether or not a piece of

code is allowed to run, and what resources it can use when it is running. For example, it is CAS that will prevent a .NET web applet from formatting your hard disk.

How does CAS work?

The CAS security policy revolves around two key concepts - code groups and permissions. Each .NET assembly is a member of a particular **code group**, and each code group is granted the permissions specified in a **named permission set**.

For example, using the default security policy, a control downloaded from a web site belongs to the 'Zone - Internet' code group, which adheres to the permissions defined by the 'Internet' named permission set. (Naturally the 'Internet' named permission set represents a very restrictive range of permissions.)

Who defines the CAS code groups?

Microsoft defines some default ones, but you can modify these and even create your own. To see the code groups defined on your system, run 'caspol -lg' from the command-line. On my system it looks like this:

- 26. Level = Machine
- 27. Code Groups:
- 28.
- 29. 1. All code: Nothing
- 30. 1.1. Zone - MyComputer: FullTrust
- 31. 1.1.1. Honor SkipVerification requests: SkipVerification
- 32. 1.2. Zone - Intranet: LocalIntranet
- 33. 1.3. Zone - Internet: Internet
- 34. 1.4. Zone - Untrusted: Nothing
- 35. 1.5. Zone - Trusted: Internet
- 1.6. StrongName -

Note the hierarchy of code groups - the top of the hierarchy is the most general ('All code'), which is then sub-divided into several groups, each of which in turn can be sub-divided. Also note that (somewhat counter-intuitively) a sub-group can be associated with a more permissive permission set than its parent.

How do I define my own code group?

Use caspol. For example, suppose you trust code from www.mydomain.com and you want it have full access to your system, but you want to keep the default restrictions for all other internet sites. To achieve this, you would add a new code group as a sub-group of the 'Zone - Internet' group, like this:

```
caspol -ag 1.3 -site www.mydomain.com FullTrust
```

Now if you run caspol -lg you will see that the new group has been added as group

1.3.1:

- 1.3. Zone - Internet: Internet
- 1.3.1. Site - www.mydomain.com: FullTrust

...

Note that the numeric label (1.3.1) is just a caspol invention to make the code groups easy to manipulate from the command-line. The underlying runtime never sees it.

How do I change the permission set for a code group?

Use caspol. If you are the machine administrator, you can operate at the 'machine' level - which means not only that the changes you make become the default for the machine, but also that users cannot change the permissions to be more permissive. If you are a normal (non-admin) user you can still modify the permissions, but only to make them more restrictive. For example, to allow intranet code to do what it likes you might do this:

```
caspol -cg 1.2 FullTrust
```

Note that because this is more permissive than the default policy (on a standard system), you should only do this at the machine level - doing it at the user level will have no effect.

Can I create my own permission set?

Yes. Use caspol -ap, specifying an XML file containing the permissions in the permission set. To save you some time, [here](#) is a sample file corresponding to the 'Everything'

permission set - just edit to suit your needs. When you have edited the sample, add it to the range of available permission sets like this:

```
caspol -ap samplepermset.xml
```

Then, to apply the permission set to a code group, do something like this:

```
caspol -cg 1.3 SamplePermSet (By default, 1.3 is the 'Internet' code group)
```

I'm having some trouble with CAS. How can I diagnose my problem?

Caspol has a couple of options that might help. First, you can ask caspol to tell you what code group an assembly belongs to, using `caspol -rsg`. Similarly, you can ask what permissions are being applied to a particular assembly using `caspol -rsp`.

I can't be bothered with all this CAS stuff. Can I turn it off?

Yes, as long as you are an administrator. Just run:

```
caspol -s off
```

http://www.codeproject.com/dotnet/UB_CAS_NET.asp

Which namespace is the base class for .net Class library?

Ans: `system.object`

What are object pooling and connection pooling and difference? Where do we set the Min and Max Pool size for connection pooling?

Object pooling is a COM+ service that enables you to reduce the overhead of creating each object from scratch. When an object is activated, it is pulled from the pool. When the object is deactivated, it is placed back into the pool to await the next request. You can configure object pooling by applying the `ObjectPoolingAttribute` attribute to a class that derives from the `System.EnterpriseServices.ServicedComponent` class.

Object pooling lets you control the number of connections you use, as opposed to connection pooling, where you control the maximum number reached.

Following are important differences between object pooling and connection pooling:

- **Creation.** When using connection pooling, creation is on the same thread, so if there is nothing in the pool, a connection is created on your behalf. With object pooling, the pool might decide to create a new object. However, if you have already reached your maximum, it instead gives you the next available object. This is crucial behavior when it takes a long time to create an object, but you do not use it for very long.
- **Enforcement of minimums and maximums.** This is not done in connection pooling. The maximum value in object pooling is very important when trying to scale your application. You might need to multiplex thousands of requests to just a few objects. (TPC/C benchmarks rely on this.)

COM+ object pooling is identical to what is used in .NET Framework managed SQL Client connection pooling. For example, creation is on a different thread and minimums and maximums are enforced.

What is Application Domain?

The primary purpose of the `AppDomain` is to isolate an application from other applications. Win32 processes provide isolation by having distinct memory address spaces. This is effective, but it is expensive and doesn't scale well. The .NET runtime enforces `AppDomain` isolation by keeping control over the use of memory - all memory in the `AppDomain` is managed by the .NET runtime, so the runtime can ensure that `AppDomains` do not access each other's memory.

Objects in different application domains communicate either by transporting copies of objects across application domain boundaries, or by using a proxy to exchange messages.

MarshalByRefObject is the base class for objects that communicate across application domain boundaries by exchanging messages using a proxy. Objects that do not inherit from **MarshalByRefObject** are implicitly marshal by value. When a remote application references a marshal by value object, a copy of the object is passed across application domain boundaries.

How does an AppDomain get created?

AppDomains are usually created by *hosts*. Examples of hosts are the Windows Shell, ASP.NET and IE. When you run a .NET application from the command-line, the host is the Shell. The Shell creates a new AppDomain for every application.

AppDomains can also be explicitly created by .NET applications. Here is a C# sample which creates an AppDomain, creates an instance of an object inside it, and then executes one of the object's methods. Note that you must name the executable 'appdomaintest.exe' for this code to work as-is.

```
using System;
using System.Runtime.Remoting;

public class CAppDomainInfo : MarshalByRefObject
{
    public string GetAppDomainInfo()
    {
        return "AppDomain = " + AppDomain.CurrentDomain.FriendlyName;
    }
}

public class App
{
    public static int Main()
    {
        AppDomain ad = AppDomain.CreateDomain( "Andy's new domain", null, null
);
        ObjectHandle oh = ad.CreateInstance( "appdomaintest",
"CAppDomainInfo" );
        CAppDomainInfo adInfo = (CAppDomainInfo)(oh.Unwrap());
        string info = adInfo.GetAppDomainInfo();
        Console.WriteLine( "AppDomain info: " + info );
        return 0;
    }
}
```

What is serialization in .NET? What are the ways to control serialization?

Serialization is the process of converting an object into a stream of bytes.

Deserialization is the opposite process of creating an object from a stream of bytes.

Serialization/Deserialization is mostly used to transport objects (e.g. during remoting), or to persist objects (e.g. to a file or database). Serialization can be defined as the process of storing the state of an object to a storage medium. During this process, the public and private fields of the object and the name of the class, including the assembly containing the class, are converted to a stream of bytes, which is then written to a data stream. When the object is subsequently deserialized, an exact clone of the original object is created.

- Binary serialization preserves type fidelity, which is useful for preserving the state of an object between different invocations of an application. For example, you can share an object between different applications by serializing it to the clipboard. You can serialize an object to a stream, disk, memory, over the network, and so forth. Remoting uses serialization to pass objects "by value" from one computer or application domain to another.
- XML serialization serializes only public properties and fields and does not preserve type fidelity. This is useful when you want to provide or consume data without restricting the application that uses the data. Because XML is an open standard, it is an attractive choice for sharing data across the Web. SOAP is an open standard, which makes it an attractive choice.

There are two separate mechanisms provided by the .NET class library - XmlSerializer and SoapFormatter/BinaryFormatter. Microsoft uses XmlSerializer for Web Services, and

uses SoapFormatter/BinaryFormatter for remoting. Both are available for use in your own code.

Why do I get errors when I try to serialize a Hashtable?

XmlSerializer will refuse to serialize instances of any class that implements IDictionary, e.g. Hashtable. SoapFormatter and BinaryFormatter do not have this restriction.

What is exception handling?

When an exception occurs, the system searches for the nearest catch clause that can handle the exception, as determined by the run-time type of the exception. First, the current method is searched for a lexically enclosing try statement, and the associated catch clauses of the try statement are considered in order. If that fails, the method that called the current method is searched for a lexically enclosing try statement that encloses the point of the call to the current method. This search continues until a catch clause is found that can handle the current exception, by naming an exception class that is of the same class, or a base class, of the run-time type of the exception being thrown. A catch clause that doesn't name an exception class can handle any exception. Once a matching catch clause is found, the system prepares to transfer control to the first statement of the catch clause. Before execution of the catch clause begins, the system first executes, in order, any finally clauses that were associated with try statements more nested than the one that caught the exception.

Exceptions that occur during destructor execution are worth special mention. If an exception occurs during destructor execution, and that exception is not caught, then the execution of that destructor is terminated and the destructor of the base class (if any) is called. If there is no base class (as in the case of the object type) or if there is no base class destructor, then the exception is discarded.

What is Assembly?

Assemblies are the building blocks of .NET Framework applications; they form the fundamental unit of deployment, version control, reuse, activation scoping, and security permissions. An assembly is a collection of types and resources that are built to work together and form a logical unit of functionality. An assembly provides the common language runtime with the information it needs to be aware of type implementations. To the runtime, a type does not exist outside the context of an assembly.

Assemblies are a fundamental part of programming with the .NET Framework. An assembly performs the following functions:

- It contains code that the common language runtime executes. Microsoft intermediate language (MSIL) code in a portable executable (PE) file will not be executed if it does not have an associated assembly manifest. Note that each assembly can have only one entry point (that is, **DllMain**, **WinMain**, or **Main**).
- It forms a security boundary. An assembly is the unit at which permissions are requested and granted.
- It forms a type boundary. Every type's identity includes the name of the assembly in which it resides. A type called MyType loaded in the scope of one assembly is not the same as a type called MyType loaded in the scope of another assembly.
- It forms a reference scope boundary. The assembly's manifest contains assembly metadata that is used for resolving types and satisfying resource requests. It specifies the types and resources that are exposed outside the assembly. The manifest also enumerates other assemblies on which it depends.
- It forms a version boundary. The assembly is the smallest versionable unit in the common language runtime; all types and resources in the same assembly are versioned as a unit. The assembly's manifest describes the version dependencies you specify for any dependent assemblies.
- It forms a deployment unit. When an application starts, only the assemblies that the application initially calls must be present. Other assemblies, such as localization resources or assemblies containing utility classes, can be retrieved on demand. This allows applications to be kept simple and thin when first downloaded.
- It is the unit at which side-by-side execution is supported.

Assemblies can be static or dynamic. Static assemblies can include .NET Framework types (interfaces and classes), as well as resources for the assembly (bitmaps, JPEG files, resource files, and so on). Static assemblies are stored on disk in PE files. You can also use the .NET Framework to create dynamic assemblies, which are run directly from memory and are not saved to disk before execution. You can save dynamic assemblies to disk after they have executed.

There are several ways to create assemblies. You can use development tools, such as Visual Studio .NET, that you have used in the past to create .dll or .exe files. You can use tools provided in the .NET Framework SDK to create assemblies with modules created in other development environments. You can also use common language runtime APIs, such as Reflection.Emit, to create dynamic assemblies.

What are the contents of assembly?

In general, a static assembly can consist of four elements:

- The assembly manifest, which contains assembly metadata.
- Type metadata.
- Microsoft intermediate language (MSIL) code that implements the types.
- A set of resources.

What are the different types of assemblies?

Private, Public/Shared, Satellite

What is the difference between a private assembly and a shared assembly?

Location and visibility: A private assembly is normally used by a single application, and is stored in the application's directory, or a sub-directory beneath. A shared assembly is normally stored in the global assembly cache, which is a repository of assemblies maintained by the .NET runtime. Shared assemblies are usually libraries of code which many applications will find useful, e.g. the .NET framework classes.

Versioning: The runtime enforces versioning constraints only on shared assemblies, not on private assemblies.

What are Satellite Assemblies? How you will create this? How will you get the different language strings?

Satellite assemblies are often used to deploy language-specific resources for an application. These language-specific assemblies work in side-by-side execution because the application has a separate product ID for each language and installs satellite assemblies in a language-specific subdirectory for each language. When uninstalling, the application removes only the satellite assemblies associated with a given language and .NET Framework version. No core .NET Framework files are removed unless the last language for that .NET Framework version is being removed.

(For example, English and Japanese editions of the .NET Framework version 1.1 share the same core files. The Japanese .NET Framework version 1.1 adds satellite assemblies with localized resources in a \ja subdirectory. An application that supports the .NET Framework version 1.1, regardless of its language, always uses the same core runtime files.)

<http://www.ondotnet.com/lpt/a/2637>

**

How will u load dynamic assembly? How will create assemblies at run time?

**

What is Assembly manifest? what all details the assembly manifest will contain?

Every assembly, whether static or dynamic, contains a collection of data that describes how the elements in the assembly relate to each other. The assembly manifest contains this assembly metadata. An assembly manifest contains all the metadata needed to specify the assembly's version requirements and security identity, and all metadata needed to define the scope of the assembly and resolve references to resources and classes. The assembly manifest can be stored in either a PE file (an .exe or .dll) with Microsoft intermediate language (MSIL) code or in a standalone PE file that contains only assembly manifest information.

It contains Assembly name, Version number, Culture, Strong name information, List of

all files in the assembly, Type reference information, Information on referenced assemblies.

Difference between assembly manifest & metadata?

assembly manifest - An integral part of every assembly that renders the assembly self-describing. The assembly manifest contains the assembly's metadata. The manifest establishes the assembly identity, specifies the files that make up the assembly implementation, specifies the types and resources that make up the assembly, itemizes the compile-time dependencies on other assemblies, and specifies the set of permissions required for the assembly to run properly. This information is used at run time to resolve references, enforce version binding policy, and validate the integrity of loaded assemblies. The self-describing nature of assemblies also helps makes zero-impact install and XCOPY deployment feasible.

metadata - Information that describes every element managed by the common language runtime: an assembly, loadable file, type, method, and so on. This can include information required for debugging and garbage collection, as well as security attributes, marshaling data, extended class and member definitions, version binding, and other information required by the runtime.

What is Global Assembly Cache (GAC) and what is the purpose of it? (How to make an assembly to public? Steps) How more than one version of an assembly can keep in same place?

Each computer where the common language runtime is installed has a machine-wide code cache called the global assembly cache. The global assembly cache stores assemblies specifically designated to be shared by several applications on the computer. You should share assemblies by installing them into the global assembly cache only when you need to.

Steps

- Create a strong name using sn.exe tool

eg: sn -k keyPair.snk

- with in AssemblyInfo.cs add the generated file name

eg: [assembly: AssemblyKeyFile("abc.snk")]

- recompile project, then install it to GAC by either

drag & drop it to assembly folder (C:\WINDOWS\assembly OR C:\WINNT\assembly) (shfusion.dll tool)

or

gacutil -i abc.dll

If I have more than one version of one assemblies, then how'll I use old version (how/where to specify version number?)in my application?

**

How to find methods of a assembly file (not using ILDASM)

Reflection

What is Garbage Collection in .Net? Garbage collection process?

The process of transitively tracing through all pointers to actively used objects in order to locate all objects that can be referenced, and then arranging to reuse any heap memory that was not found during this trace. The common language runtime garbage collector also compacts the memory that is in use to reduce the working space needed for the heap.

What is Reflection in .NET? Namespace? How will you load an assembly which is not referenced by current assembly?

All .NET compilers produce metadata about the types defined in the modules they produce. This metadata is packaged along with the module (modules in turn are packaged together in assemblies), and can be accessed by a mechanism called **reflection**. The System.Reflection namespace contains classes that can be used to interrogate the types for a module/assembly.

Using reflection to access .NET metadata is very similar to using ITypeLib/ITypeInfo to access type library data in COM, and it is used for similar purposes - e.g. determining data type sizes for marshaling data across context/process/machine boundaries.

Reflection can also be used to dynamically invoke methods (see

System.Type.InvokeMember), or even create types dynamically at run-time (see System.Reflection.Emit.TypeBuilder).

What is Custom attribute? How to create? If I'm having custom attribute in an assembly, how to say that name in the code?

A: The primary steps to properly design custom attribute classes are as follows:

```
Applying the AttributeUsageAttribute ([AttributeUsage(AttributeTargets.All,
Inherited = false, AllowMultiple = true)])
Declaring the attribute. (class public class MyAttribute : System.Attribute { // . .
. })
Declaring constructors (public MyAttribute(bool myvalue) { this.myvalue =
myvalue; })
Declaring properties
public bool MyProperty
{
get {return this.myvalue;}
set {this.myvalue = value;}
}
```

The following example demonstrates the basic way of using reflection to get access to custom attributes.

```
class MainClass
{
public static void Main()
{
System.Reflection.MemberInfo info = typeof(MyClass);
object[] attributes = info.GetCustomAttributes();
for (int i = 0; i < attributes.Length; i++)
{
System.Console.WriteLine(attributes[i]);
}
}
}
```

What is the managed and unmanaged code in .net?

The .NET Framework provides a run-time environment called the Common Language Runtime, which manages the execution of code and provides services that make the development process easier. Compilers and tools expose the runtime's functionality and enable you to write code that benefits from this managed execution environment. Code that you develop with a language compiler that targets the runtime is called *managed code*; it benefits from features such as cross-language integration, cross-language exception handling, enhanced security, versioning and deployment support, a simplified model for component interaction, and debugging and profiling services.

How do you create threading in .NET? What is the namespace for that?

**

System.Threading.Thread

Serialize and MarshalByRef?

using directive vs **using** statement

You create an instance in a **using** statement to ensure that **Dispose** is called on the object when the **using** statement is exited. A **using** statement can be exited either when the end of the **using** statement is reached or if, for example, an exception is thrown and control leaves the statement block before the end of the statement.

The **using** directive has two uses:

- Create an alias for a namespace (a **using** alias).
- Permit the use of types in a namespace, such that, you do not have to qualify the use of a type in that namespace (a **using** directive).

Describe the Managed Execution Process?

The managed execution process includes the following steps:

1. Choosing a compiler.
To obtain the benefits provided by the common language runtime, you must use one or more language compilers that target the runtime.
2. Compiling your code to Microsoft intermediate language (MSIL).
Compiling translates your source code into MSIL and generates the required metadata.
3. Compiling MSIL to native code.
At execution time, a just-in-time (JIT) compiler translates the MSIL into native code. During this compilation, code must pass a verification process that examines the MSIL and metadata to find out whether the code can be determined to be type safe.
4. Executing your code.
The common language runtime provides the infrastructure that enables execution to take place as well as a variety of services that can be used during execution.

What is Active Directory? What is the namespace used to access the Microsoft Active Directories? What are ADSI Directories?

Active Directory Service Interfaces (ADSI) is a programmatic interface for Microsoft Windows Active Directory. It enables your applications to interact with diverse directories on a network, using a single interface. Visual Studio .NET and the .NET Framework make it easy to add ADSI functionality with the **DirectoryEntry** and **DirectorySearcher** components.

Using ADSI, you can create applications that perform common administrative tasks, such as backing up databases, accessing printers, and administering user accounts. ADSI makes it possible for you to:

- Log on once to work with diverse directories. The **DirectoryEntry** component class provides username and password properties that can be entered at runtime and communicated to the Active Directory object you are binding to.
- Use a single application programming interface (API) to perform tasks on multiple directory systems by offering the user a variety of protocols to use. The **DirectoryServices** namespace provides the classes to perform most administrative functions.
- Perform "rich querying" on directory systems. ADSI technology allows for searching for an object by specifying two query dialects: SQL and LDAP.
- Access and use a single, hierarchical structure for administering and maintaining diverse and complicated network configurations by accessing an Active Directory tree.
- Integrate directory information with databases such as SQL Server. The **DirectoryEntry** path may be used as an ADO.NET connection string provided that it is using the LDAP provider.

using System.DirectoryServices;

How Garbage Collector (GC) Works?

The methods in this class influence when an object is garbage collected and when resources allocated by an object are released. Properties in this class provide information about the total amount of memory available in the system and the age category, or generation, of memory allocated to an object. Periodically, the garbage collector performs garbage collection to reclaim memory allocated to objects for which there are no valid references. Garbage collection happens automatically when a request for memory cannot be satisfied using available free memory. Alternatively, an application can force garbage collection using the Collect method.

Garbage collection consists of the following steps:

1. The garbage collector searches for managed objects that are referenced in managed code.
2. The garbage collector attempts to finalize objects that are not referenced.

3. The garbage collector frees objects that are not referenced and reclaims their memory.

Why do we need to call CG.SuppressFinalize?

Requests that the system not call the finalizer method for the specified object.

public static void SuppressFinalize(

object *obj*

); The method removes *obj* from the set of objects that require finalization. The *obj* parameter is required to be the caller of this method.

Objects that implement the IDisposable interface can call this method from the IDisposable.Dispose method to prevent the garbage collector from calling Object.Finalize on an object that does not require it.

What is nmake tool?

The Nmake tool (Nmake.exe) is a 32-bit tool that you use to build projects based on commands contained in a .mak file.

usage : **nmake -a all**

What are Namespaces?

The **namespace** keyword is used to declare a scope. This namespace scope lets you organize code and gives you a way to create globally-unique types. Even if you do not explicitly declare one, a default namespace is created. This unnamed namespace, sometimes called the global namespace, is present in every file. Any identifier in the global namespace is available for use in a named namespace. Namespaces implicitly have public access and this is not modifiable.

What is the difference between CONST and READONLY?

Both are meant for constant values. A **const** field can only be initialized at the declaration of the field. A **readonly** field can be initialized either at the declaration or in a constructor. Therefore, **readonly** fields can have different values depending on the constructor used.

```
readonly int b;  
public X()  
{  
b=1;  
}  
public X(string s)  
{  
b=5;  
}  
public X(string s, int i)  
{  
b=i;  
}
```

Also, while a **const** field is a compile-time constant, the **readonly** field can be used for runtime constants, as in the following example:

```
public static readonly uint l1 = (uint) DateTime.Now.Ticks; (this can't be possible with const)
```

What is the difference between ref & out parameters?

An argument passed to a ref parameter must first be initialized. Compare this to an out parameter, whose argument does not have to be explicitly initialized before being passed to an out parameter.

What is the difference between Array and LinkedList?

What is the difference between Array and ArrayList?

As elements are added to an ArrayList, the capacity is automatically increased as required through reallocation. The capacity can be decreased by calling TrimToSize or by setting the Capacity property explicitly.

What is Jagged Arrays?

A jagged array is an array whose elements are arrays. The elements of a jagged array can be of different dimensions and sizes. A jagged array is sometimes called an "array-of-arrays."

What are indexers?

Indexers are similar to properties, except that the **get** and **set** accessors of indexers take parameters, while property accessors do not.

What is Asynchronous call and how it can be implemented using delegates?

How to create events for a control? What is custom events? How to create it?

If you want to write your own dot net language, what steps you will u take care?

61. Describe the difference between inline and code behind - which is best in a loosely coupled solution?

how dot net compiled code will become platform independent?

without modifying source code if we compile again, will it be generated MSIL again?

C++ & C# differences

**

(COM)

Interop Services?

The common language runtime provides two mechanisms for interoperating with unmanaged code:

- Platform invoke, which enables managed code to call functions exported from an unmanaged library.
- COM interop, which enables managed code to interact with COM objects through interfaces.

Both platform invoke and COM interop use interop marshaling to accurately move method arguments between caller and callee and back, if required.

How does u handle this COM components developed in other programming languages in .NET?

What is RCW (Runtime Callable Wrappers)?

The common language runtime exposes COM objects through a proxy called the runtime callable wrapper (RCW). Although the RCW appears to be an ordinary object to .NET clients, its primary function is to marshal calls between a .NET client and a COM object.

What is CCW (COM Callable Wrapper)

A proxy object generated by the common language runtime so that existing COM applications can use managed classes, including .NET Framework classes, transparently.

How CCW and RCW is working?

**

How will you register com+ services?

The .NET Framework SDK provides the .NET Framework Services Installation Tool (Regsvcs.exe - a command-line tool) to manually register an assembly containing serviced components. You can also access these registration features programmatically with the System.EnterpriseServices.RegistrationHelper class by creating an instance of class RegistrationHelper and using the method InstallAssembly

What is use of ContextUtil class?

ContextUtil is the preferred class to use for obtaining COM+ context information.

What is the new three features of COM+ services, which are not there in COM (MTS)?

**

Is the COM architecture same as .Net architecture? What is the difference between them?

**

Can we copy a COM dll to GAC folder?

**

What is Pinvoke?

Platform invoke is a service that enables managed code to call unmanaged functions implemented in dynamic-link libraries (DLLs), such as those in the Win32 API. It locates and invokes an exported function and marshals its arguments (integers, strings, arrays, structures, and so on) across the interoperation boundary as needed.

Is it true that COM objects no longer need to be registered on the server?

Answer: Yes and No. Legacy COM objects still need to be registered on the server before they can be used. COM developed using the new .NET Framework will not need to be registered. Developers will be able to auto-register these objects just by placing them in the 'bin' folder of the application.

Can .NET Framework components use the features of Component Services?

Answer: Yes, you can use the features and functions of Component Services from a .NET Framework component.

<http://msdn.microsoft.com/library/techart/Pahlcompserv.htm>

(OOPS)

What are the OOPS concepts?

- 1) Encapsulation: It is the mechanism that binds together code and data in manipulates, and keeps both safe from outside interference and misuse. In short it isolates a particular code and data from all other codes and data. A well-defined interface controls the access to that particular code and data.
- 2) Inheritance: It is the process by which one object acquires the properties of another object. This supports the hierarchical classification. Without the use of hierarchies, each object would need to define all its characteristics explicitly. However, by use of inheritance, an object need only define those qualities that make it unique within its class. It can inherit its general attributes from its parent. A new sub-class inherits all of the attributes of all of its ancestors.
- 3) Polymorphism: It is a feature that allows one interface to be used for general class of actions. The specific action is determined by the exact nature of the situation. In general polymorphism means "one interface, multiple methods", This means that it is possible to design a generic interface to a group of related activities. This helps reduce complexity by allowing the same interface to be used to specify a general class of action. It is the compiler's job to select the specific action (that is, method) as it applies to each situation.

What is the difference between a Struct and a Class?

- The struct type is suitable for representing lightweight objects such as Point, Rectangle, and Color. Although it is possible to represent a point as a class, a struct is more efficient in some scenarios. For example, if you declare an array of 1000 Point objects, you will allocate additional memory for referencing each object. In this case, the struct is less expensive.
- When you create a struct object using the new operator, it gets created and the appropriate constructor is called. Unlike classes, structs can be instantiated without using the new operator. If you do not use new, the fields will remain unassigned and the object cannot be used until all of the fields are initialized.
- It is an error to declare a default (parameterless) constructor for a struct. A default constructor is always provided to initialize the struct members to their default values.
- It is an error to initialize an instance field in a struct.
- There is no inheritance for structs as there is for classes. A struct cannot inherit from another struct or class, and it cannot be the base of a class. Structs, however, inherit from the base class Object. A struct can implement interfaces, and it does that exactly as classes do.
- A struct is a value type, while a class is a reference type.

Value type & reference types difference? Example from .NET. Integer & struct are value types or reference types in .NET?

Most programming languages provide built-in data types, such as integers and floating-point numbers, that are copied when they are passed as arguments (that is, they are

passed by value). In the .NET Framework, these are called value types. The runtime supports two kinds of value types:

- Built-in value types
The .NET Framework defines built-in value types, such as System.Int32 and System.Boolean, which correspond and are identical to primitive data types used by programming languages.
- User-defined value types
Your language will provide ways to define your own value types, which derive from System.ValueType. If you want to define a type representing a value that is small, such as a complex number (using two floating-point numbers), you might choose to define it as a value type because you can pass the value type efficiently by value. If the type you are defining would be more efficiently passed by reference, you should define it as a class instead.

Variables of reference types, referred to as objects, store references to the actual data. The following are the reference types:

- class
- interface
- delegate

The following are the built-in reference types:

- object
- string

What is Inheritance, Multiple Inheritance, Shared and Repeatable Inheritance?

**

What is Method overloading?

Method overloading occurs when a class contains two methods with the same name, but different signatures.

What is Method Overriding? How to override a function in C#?

Use the override modifier to modify a method, a property, an indexer, or an event. An override method provides a new implementation of a member inherited from a base class. The method overridden by an override declaration is known as the overridden base method. The overridden base method must have the same signature as the override method.

You cannot override a non-virtual or static method. The overridden base method must be virtual, abstract, or override.

Can we call a base class method without creating instance?

It's possible if it's a static method.

It's possible by inheriting from that class also.

It's possible from derived classes using the base keyword.

You have one base class virtual function how will call that function from derived class?

Ans:

```
22. class a
23.     {
24.         public virtual int m()
25.         {
26.             return 1;
27.         }
28.     }
29.     class b:a
30.     {
31.         public int j()
32.         {
33.             return m();
```

```
34.     }
        }
```

In which cases you use override and new base?

Use the new modifier to explicitly hide a member inherited from a base class. To hide an inherited member, declare it in the derived class using the same name, and modify it with the new modifier.

C# Language features

What are Sealed Classes in C#?

The sealed modifier is used to prevent derivation from a class. A compile-time error occurs if a sealed class is specified as the base class of another class. (A sealed class cannot also be an abstract class)

What is Polymorphism? How does VB.NET/C# achieve polymorphism?

**

```
38. class Token
39.     {
40.         public string Display()
41.         {
42.             //Implementation goes here
43.             return "base";
44.         }
45.     }
46.     class IdentifierToken:Token
47.     {
48.         public new string Display() //What is the use of new keyword
49.         {
50.             //Implementation goes here
51.             return "derive";
52.         }
53.     }
54.     static void Method(Token t)
55.     {
56.         Console.Write(t.Display());
57.     }
58.     public static void Main()
59.     {
60.         IdentifierToken Variable=new IdentifierToken();
61.         Method(Variable); //Which Class Method is called here
62.         Console.ReadLine();
63.     }
64.     For the above code What is the "new" keyword and Which Class Method is
65.     called here
```

A: it will call base class Display method

```
66. class Token
67.     {
68.         public virtual string Display()
69.         {
70.             //Implementation goes here
71.             return "base";
72.         }
73.     }
74.     class IdentifierToken:Token
75.     {
76.         public override string Display() //What is the use of new keyword
77.         {
```

```

78.         //Implementation goes here
79.         return "derive";
80.     }
81. }
82. static void Method(Token t)
83. {
84.     Console.Write(t.Display());
85. }
86. public static void Main()
87. {
88.     IdentifierToken Variable=new IdentifierToken();
89.     Method(Variable); //Which Class Method is called here
90.     Console.ReadLine();
91. }
92. A: Derive

```

In which Scenario you will go for Interface or Abstract Class?

Interfaces, like classes, define a set of properties, methods, and events. But unlike classes, interfaces do not provide implementation. They are implemented by classes, and defined as separate entities from classes. Even though class inheritance allows your classes to inherit implementation from a base class, it also forces you to make most of your design decisions when the class is first published.

Abstract classes are useful when creating components because they allow you specify an invariant level of functionality in some methods, but leave the implementation of other methods until a specific implementation of that class is needed. They also version well, because if additional functionality is needed in derived classes, it can be added to the base class without breaking code.

Interfaces vs. Abstract Classes

Feature	Interface	Abstract class
Multiple inheritance	A class may implement several interfaces.	A class may extend only one abstract class.
Default implementation	An interface cannot provide any code at all, much less default code.	An abstract class can provide complete code, default code, and/or just stubs that have to be overridden.
Constants	Static final constants only, can use them without qualification in classes that implement the interface. On the other paw, these unqualified names pollute the namespace. You can use them and it is not obvious where they are coming from since the qualification is optional.	Both instance and static constants are possible. Both static and instance intialiser code are also possible to compute the constants.
Third party convenience	An interface implementation may be added to any existing third party class.	A third party class must be rewritten to extend only from the abstract class.
is-a vs -able or can-do	Interfaces are often used to describe the peripheral abilities of a class, not its central identity, e.g. an Automobile class might implement the Recyclable interface, which could apply to many otherwise	An abstract class defines the core identity of its descendants. If you defined a Dog abstract class then Damamation descendants are Dogs, they are not merely dogable. Implemented interfaces enumerate the general things a class can do, not the things a

	totally unrelated objects.	class is.
Plug-in	You can write a new replacement module for an interface that contains not one stick of code in common with the existing implementations. When you implement the interface, you start from scratch without any default implementation. You have to obtain your tools from other classes; nothing comes with the interface other than a few constants. This gives you freedom to implement a radically different internal design.	You must use the abstract class as-is for the code base, with all its attendant baggage, good or bad. The abstract class author has imposed structure on you. Depending on the cleverness of the author of the abstract class, this may be good or bad. Another issue that's important is what I call "heterogeneous vs. homogeneous." If implementors/subclasses are homogeneous, tend towards an abstract base class. If they are heterogeneous, use an interface. (Now all I have to do is come up with a good definition of hetero/homogeneous in this context.) If the various objects are all of-a-kind, and share a common state and behavior, then tend towards a common base class. If all they share is a set of method signatures, then tend towards an interface.
Homogeneity	If all the various implementations share is the method signatures, then an interface works best.	If the various implementations are all of a kind and share a common status and behavior, usually an abstract class works best.
Maintenance	If your client code talks only in terms of an interface, you can easily change the concrete implementation behind it, using a factory method.	Just like an interface, if your client code talks only in terms of an abstract class, you can easily change the concrete implementation behind it, using a factory method.
Speed	Slow, requires extra indirection to find the corresponding method in the actual class. Modern JVMs are discovering ways to reduce this speed penalty.	Fast
Terseness	The constant declarations in an interface are all presumed public static final, so you may leave that part out. You can't call any methods to compute the initial values of your constants. You need not declare individual methods of an interface abstract. They are all presumed so.	You can put shared code into an abstract class, where you cannot into an interface. If interfaces want to share code, you will have to write other bubblegum to arrange that. You may use methods to compute the initial values of your constants and variables, both instance and static. You must declare all the individual methods of an abstract class abstract.
Adding	If you add a new method	If you add a new method to an

functionality	to an interface, you must track down all implementations of that interface in the universe and provide them with a concrete implementation of that method.	abstract class, you have the option of providing a default implementation of it. Then all existing code will continue to work without change.
---------------	--	---

94. see the code

```

95. interface ICommon
96.     {
97.         int getCommon();
98.     }
99.     interface ICommonImplements1:ICommon
100.    {
101.    }
102.    interface ICommonImplements2:ICommon
103.    {
104.    }
105.    public class a:ICommonImplements1,ICommonImplements2
106.    {
    }

```

How to implement getCommon method in class a? Are you seeing any problem in the implementation?

Ans:

```

public class a:ICommonImplements1,ICommonImplements2
    {
        public int getCommon()
        {
            return 1;
        }
    }
107. interface IWeather
108.     {
109.         void display();
110.     }
111.     public class A:IWeather
112.     {
113.         public void display()
114.         {
115.             MessageBox.Show("A");
116.         }
117.     }
118.     public class B:A
119.     {
120.     }
121.     public class C:B,IWeather
122.     {
123.         public void display()
124.         {
125.             MessageBox.Show("C");
126.         }
127.     }
128. When I instantiate C.display(), will it work?
129. interface IPrint
130.     {
131.         string Display();

```

```

132.     }
133.     interface IWrite
134.     {
135.         string Display();
136.     }
137.     class PrintDoc:IPrint,IWrite
138.     {
139.         //Here is implementation
140.     }

```

how to implement the Display in the class printDoc (How to resolve the naming Conflict) A: no naming conflicts

```

class PrintDoc:IPrint,IWrite
{
    public string Display()
    {
        return "s";
    }
}
141.interface IList
142.     {
143.         int Count { get; set; }
144.     }
145.     interface ICounter
146.     {
147.         void Count(int i);
148.     }
149.     interface IListCounter: IList, ICounter {}
150.     class C
151.     {
152.         void Test(IListCounter x)
153.         {
154.             x.Count(1);           // Error
155.             x.Count = 1;         // Error
156.             ((IList)x).Count = 1; // Ok, invokes IList.Count.set
157.             ((ICounter)x).Count(1); // Ok, invokes ICounter.Count
158.         }
159.     }

```

Write one code example for compile time binding and one for run time binding? What is early/late binding?

An object is *early bound* when it is assigned to a variable declared to be of a specific object type. Early bound objects allow the compiler to allocate memory and perform other optimizations before an application executes.

' Create a variable to hold a new object.

```
Dim FS As FileStream
```

' Assign a new object to the variable.

```
FS = New FileStream("C:\tmp.txt", FileMode.Open)
```

By contrast, an object is *late bound* when it is assigned to a variable declared to be of type **Object**. Objects of this type can hold references to any object, but lack many of the advantages of early-bound objects.

```
Dim xlApp As Object
```

```
xlApp = CreateObject("Excel.Application")
```

161.Can you explain what inheritance is and an example of when you might use it?

How can you write a class to restrict that only one object of this class can be created (Singleton class)?

(Access specifiers)

What are the access-specifiers available in c#?

Private, Protected, Public, Internal, Protected Internal.

Explain about Protected and protected internal, "internal" access-specifier?

protected - Access is limited to the containing class or types derived from the containing class.

internal - Access is limited to the current assembly.

protected internal - Access is limited to the current assembly or types derived from the containing class.

(Constructor / Destructor)

Difference between type constructor and instance constructor? What is static constructor, when it will be fired? And what is its use?

(Class constructor method is also known as type constructor or type initializer)

Instance constructor is executed when a new instance of type is created and the class constructor is executed after the type is loaded and before any one of the type members is accessed. (It will get executed only 1st time, when we call any static methods/fields in the same class.) Class constructors are used for static field initialization. Only one class constructor per type is permitted, and it cannot use the vararg (variable argument) calling convention.

A static constructor is used to initialize a class. It is called automatically to initialize the class before the first instance is created or any static members are referenced.

What is Private Constructor? and it's use? Can you create instance of a class which has Private Constructor?

A: When a class declares only private instance constructors, it is not possible for classes outside the program to derive from the class or to directly create instances of it. (Except Nested classes)

Make a constructor private if:

- You want it to be available only to the class itself. For example, you might have a special constructor used only in the implementation of your class' Clone method.
- You do not want instances of your component to be created. For example, you may have a class containing nothing but Shared utility functions, and no instance data. Creating instances of the class would waste memory.

I have 3 overloaded constructors in my class. In order to avoid making instance of the class do I need to make all constructors to private?

(yes)

Overloaded constructor will call default constructor internally?

(no)

What are virtual destructors?

Destructor and finalize

Generally in C++ the destructor is called when objects gets destroyed. And one can explicitly call the destructors in C++. And also the objects are destroyed in reverse order that they are created in. So in C++ you have control over the destructors.

In C# you can never call them, the reason is one cannot destroy an object. So who has the control over the destructor (in C#)? it's the .Net frameworks Garbage Collector (GC). GC destroys the objects only when necessary. Some situations of necessity are memory is exhausted or user explicitly calls System.GC.Collect() method.

Points to remember:

1. Destructors are invoked automatically, and cannot be invoked explicitly.
2. Destructors cannot be overloaded. Thus, a class can have, at most, one destructor.
3. Destructors are not inherited. Thus, a class has no destructors other than the one, which may be declared in it.
4. Destructors cannot be used with structs. They are only used with classes.
5. An instance becomes eligible for destruction when it is no longer possible for any code to use the instance.
6. Execution of the destructor for the instance may occur at any time after the instance becomes eligible for destruction.
7. When an instance is destructed, the destructors in its inheritance chain are called, in order, from most derived to least derived.

<http://msdn.microsoft.com/library/default.asp?url=/library/en-us/cpguide/html/cpconfinalizemethodscdestructors.asp>

What is the difference between Finalize and Dispose (Garbage collection)

Class instances often encapsulate control over resources that are not managed by the runtime, such as window handles (HWND), database connections, and so on.

Therefore, you should provide both an explicit and an implicit way to free those resources. Provide implicit control by implementing the protected Finalize Method on an object (destructor syntax in C# and the Managed Extensions for C++). The garbage collector calls this method at some point after there are no longer any valid references to the object.

In some cases, you might want to provide programmers using an object with the ability to explicitly release these external resources before the garbage collector frees the object. If an external resource is scarce or expensive, better performance can be achieved if the programmer explicitly releases resources when they are no longer being used. To provide explicit control, implement the Dispose method provided by the IDisposable Interface. The consumer of the object should call this method when it is done using the object. **Dispose** can be called even if other references to the object are alive.

Note that even when you provide explicit control by way of **Dispose**, you should provide implicit cleanup using the **Finalize** method. **Finalize** provides a backup to prevent resources from permanently leaking if the programmer fails to call **Dispose**.

What is close method? How its different from Finalize & Dispose?

**

What is boxing & unboxing?

What is check/uncheck?

What is the use of base keyword? Tell me a practical example for base keyword's usage?

176. What are the different .net tools which u used in projects?

```
try
{
...
}
catch
{
...//exception occurred here. What'll happen?
}
finally
{
..
}
```

Ans : It will throw exception.

178. What will do to avoid prior case?

Ans:

179. try

180. {

181. try

182. {

183.

184. }

185. catch

186. {

187.

188. //exception occurred here.

189. }

190. finally

191. {

192.

193. }

194. }

```

195.catch
196.{
197....
198.}
199.finally
200.{
201....
    }
202.try
203.{
204....
205.}
206.catch
207.{
208....
209.}
210.finally
211.{
212...
213.}

```

Will it go to finally block if there is no exception happened?

Ans: Yes. The **finally** block is useful for cleaning up any resources allocated in the try block. Control is always passed to the finally block regardless of how the try block exits.

Is goto statement supported in C#? How about Java?

Gotos are supported in C# to the fullest. In Java goto is a reserved keyword that provides absolutely no functionality.

What's different about switch statements in C#?

No fall-throughs allowed. Unlike the C++ **switch** statement, C# does not support an explicit fall through from one case label to another. If you want, you can use **goto** a switch-case, or **goto default**.

```

case 1:
cost += 25;
break;
case 2:
cost += 25;
goto case 1;

```

(ADO.NET)

Advantage of ADO.Net?

- ADO.NET Does Not Depend On Continuously Live Connections
- Database Interactions Are Performed Using Data Commands
- Data Can Be Cached in Datasets
- Datasets Are Independent of Data Sources
- Data Is Persisted as XML
- Schemas Define Data Structures

How would u connect to database using .NET?

```

SqlConnection nwindConn = new SqlConnection("Data Source=localhost; Integrated
Security=SSPI;" +

```

```

"Initial Catalog=northwind");

```

```

nwindConn.Open();

```

What are relation objects in dataset and how & where to use them?

In a **DataSet** that contains multiple **DataTable** objects, you can use **DataRelation** objects to relate one table to another, to navigate through the tables, and to return child or parent rows from a related table. Adding a **DataRelation** to a **DataSet** adds, by default, a **UniqueConstraint** to the parent table and a **ForeignKeyConstraint** to the child table.

The following code example creates a **DataRelation** using two **DataTable** objects in a **DataSet**. Each **DataTable** contains a column named **CustID**, which serves as a link between the two **DataTable** objects. The example adds a single **DataRelation** to the

Relations collection of the **DataSet**. The first argument in the example specifies the name of the **DataRelation** being created. The second argument sets the parent **DataColumn** and the third argument sets the child **DataColumn**.

```
custDS.Relations.Add("CustOrders",  
custDS.Tables["Customers"].Columns["CustID"],  
custDS.Tables["Orders"].Columns["CustID"]);
```

OR

```
private void CreateRelation()  
{  
// Get the DataColumn objects from two DataTable objects in a DataSet.  
DataColumn parentCol;  
DataColumn childCol;  
// Code to get the DataSet not shown here.  
parentCol = DataSet1.Tables["Customers"].Columns["CustID"];  
childCol = DataSet1.Tables["Orders"].Columns["CustID"];  
// Create DataRelation.  
DataRelation relCustOrder;  
relCustOrder = new DataRelation("CustomersOrders", parentCol, childCol);  
// Add the relation to the DataSet.  
DataSet1.Relations.Add(relCustOrder);  
}
```

Difference between OLEDB Provider and SqlClient ?

Ans: SQLClient .NET classes are highly optimized for the .net / sqlserver combination and achieve optimal results. The SqlClient data provider is fast. It's faster than the Oracle provider, and faster than accessing database via the OleDb layer. It's faster because it accesses the native library (which automatically gives you better performance), and it was written with lots of help from the SQL Server team.

What are the different namespaces used in the project to connect the database? What data providers available in .net to connect to database?

- System.Data.OleDb – classes that make up the .NET Framework Data Provider for OLE DB-compatible data sources. These classes allow you to connect to an OLE DB data source, execute commands against the source, and read the results.
- System.Data.SqlClient – classes that make up the .NET Framework Data Provider for SQL Server, which allows you to connect to SQL Server 7.0, execute commands, and read results. The **System.Data.SqlClient** namespace is similar to the **System.Data.OleDb** namespace, but is optimized for access to SQL Server 7.0 and later.
- System.Data.Odbc - classes that make up the .NET Framework Data Provider for ODBC. These classes allow you to access ODBC data source in the managed space.
- System.Data.OracleClient - classes that make up the .NET Framework Data Provider for Oracle. These classes allow you to access an Oracle data source in the managed space.

Difference between DataReader and DataAdapter / DataSet and DataAdapter?

You can use the ADO.NET DataReader to retrieve a read-only, forward-only stream of data from a database. Using the DataReader can increase application performance and reduce system overhead because only one row at a time is ever in memory.

After creating an instance of the **Command** object, you create a **DataReader** by calling **Command.ExecuteReader** to retrieve rows from a data source, as shown in the following example.

```
SqlDataReader myReader = myCommand.ExecuteReader();
```

You use the **Read** method of the **DataReader** object to obtain a row from the results of the query.

```
while (myReader.Read())
```

```
    Console.WriteLine("\t{0}\t{1}", myReader.GetInt32(0), myReader.GetString(1));  
myReader.Close();
```

The DataSet is a memory-resident representation of data that provides a consistent relational programming model regardless of the data source. It can be used with multiple and differing data sources, used with XML data, or used to manage data local to the application. The **DataSet** represents a complete set of data including related tables, constraints, and relationships among the tables. The methods and objects in a **DataSet** are consistent with those in the relational database model. The **DataSet** can also persist and reload its contents as XML and its schema as XML Schema definition language (XSD) schema.

The DataAdapter serves as a bridge between a DataSet and a data source for retrieving and saving data. The DataAdapter provides this bridge by mapping Fill, which changes the data in the DataSet to match the data in the data source, and Update, which changes the data in the data source to match the data in the DataSet. If you are connecting to a Microsoft SQL Server database, you can increase overall performance by using the SqlDataAdapter along with its associated SqlCommand and SqlConnection. For other OLE DB-supported databases, use the DataAdapter with its associated OleDbCommand and OleDbConnection objects.

Which method do you invoke on the DataAdapter control to load your generated dataset with data?

Fill()

Explain different methods and Properties of DataReader which you have used in your project?

Read

GetString

GetInt32

```
while (myReader.Read())
```

```
    Console.WriteLine("\t{0}\t{1}", myReader.GetInt32(0), myReader.GetString(1));
```

```
myReader.Close();
```

What happens when we issue DataSet.ReadXml command?

Reads XML schema and data into the DataSet.

In how many ways we can retrieve table records count? How to find the count of records in a dataset?

```
foreach(DataTable thisTable in myDataSet.Tables){
```

```
// For each row, print the values of each column.
```

```
foreach(DataRow myRow in thisTable.Rows){
```

How to check if a datareader is closed or opened?

```
IsClosed()
```

What happens when u try to update data in a dataset in .NET while the record is already deleted in SQL SERVER as backend?

OR What is concurrency? How will you avoid concurrency when dealing with dataset? (One user deleted one row after that another user through his dataset was trying to update same row. What will happen? How will you avoid the problem?)

**

How do you merge 2 datasets into the third dataset in a simple manner? OR If you are executing these statements in commandObject. "Select * from Table1;Select * from Table2" how you will deal result set?

**

How do you sort a dataset?

**

If a dataset contains 100 rows, how to fetch rows between 5 and 15 only?

**

Differences between dataset.clone and dataset.copy?

Clone - Copies the structure of the DataSet, including all DataTable schemas, relations, and constraints. Does not copy any data.

Copy - Copies both the structure and data for this DataSet.

What is the use of parameter object?

**

How to generate XML from a dataset and vice versa?

**

What is method to get XML and schema from Dataset?

ans: getXML () and get Schema ()

How do u implement locking concept for dataset?

**

(ASP.NET)

Asp.net and asp – differences?

Code Render Block	Code Declaration Block
	Compiled
Request/Response	Event Driven
	Object Oriented - Constructors/Destructors, Inheritance, overloading..
	Exception Handling - Try, Catch, Finally
	Down-level Support
	Cultures
	User Controls
	In-built client side validation
Session - weren't transferable across servers	It can span across servers, It can survive server crashes, can work with browsers that don't support cookies
built on top of the window & IIS, it was always a separate entity & its functionality was limited.	its an integral part of OS under the .net framework. It shares many of the same objects that traditional applications would use, and all .net objects are available for asp.net's consumption.
	Garbage Collection
	Declare variable with datatype
	In built graphics support
	Cultures

How ASP and ASP.NET page works? Explain about asp.net page life cycle?

**

Order of events in an asp.net page? Control Execution Lifecycle?

Phase	What a control needs to do	Method or event to override
Initialize	Initialize settings needed during the lifetime of the incoming Web request.	Init event (OnInit method)
Load view state	At the end of this phase, the ViewState property of a control is automatically populated as described in Maintaining State in a Control. A control can override the default implementation of the LoadViewState method to customize state restoration.	LoadViewState method
Process postback data	Process incoming form data and update properties accordingly.	LoadPostData method (if IPostBackDataHandler is implemented)
Load	Perform actions common to all	Load event

	requests, such as setting up a database query. At this point, server controls in the tree are created and initialized, the state is restored, and form controls reflect client-side data.	(OnLoad method)
Send postback change notifications	Raise change events in response to state changes between the current and previous postbacks.	RaisePostDataChangedEvent method (if IPostBackDataHandler is implemented)
Handle postback events	Handle the client-side event that caused the postback and raise appropriate events on the server.	RaisePostBackEvent method (if IPostBackEventHandler is implemented)
Prerender	Perform any updates before the output is rendered. Any changes made to the state of the control in the prerender phase can be saved, while changes made in the rendering phase are lost.	PreRender event (OnPreRender method)
Save state	The ViewState property of a control is automatically persisted to a string object after this stage. This string object is sent to the client and back as a hidden variable. For improving efficiency, a control can override the SaveViewState method to modify the ViewState property.	SaveViewState method
Render	Generate output to be rendered to the client.	Render method
Dispose	Perform any final cleanup before the control is torn down. References to expensive resources such as database connections must be released in this phase.	Dispose method
Unload	Perform any final cleanup before the control is torn down. Control authors generally perform cleanup in Dispose and do not handle this event.	Unload event (On UnLoad method)

Note To override an *EventName* event, override the *OnEventName* method (and call base. *OnEventName*).

What are server controls?

ASP.NET server controls are components that run on the server and encapsulate user-interface and other related functionality. They are used in ASP.NET pages and in ASP.NET code-behind classes.

What is the difference between Web User Control and Web Custom Control? Custom Controls

Web custom controls are compiled components that run on the server and that encapsulate user-interface and other related functionality into reusable packages. They can include all the design-time features of standard ASP.NET server controls, including full support for Visual Studio design features such as the Properties window, the visual

designer, and the Toolbox.

There are several ways that you can create Web custom controls:

- You can compile a control that combines the functionality of two or more existing controls. For example, if you need a control that encapsulates a button and a text box, you can create it by compiling the existing controls together.
- If an existing server control almost meets your requirements but lacks some required features, you can customize the control by deriving from it and overriding its properties, methods, and events.
- If none of the existing Web server controls (or their combinations) meet your requirements, you can create a custom control by deriving from one of the base control classes. These classes provide all the basic functionality of Web server controls, so you can focus on programming the features you need.

If none of the existing ASP.NET server controls meet the specific requirements of your applications, you can create either a Web user control or a Web custom control that encapsulates the functionality you need. The main difference between the two controls lies in ease of creation vs. ease of use at design time.

Web user controls are easy to make, but they can be less convenient to use in advanced scenarios. You develop Web user controls almost exactly the same way that you develop Web Forms pages. Like Web Forms, user controls can be created in the visual designer, they can be written with code separated from the HTML, and they can handle execution events. However, because Web user controls are compiled dynamically at run time they cannot be added to the Toolbox, and they are represented by a simple placeholder glyph when added to a page. This makes Web user controls harder to use if you are accustomed to full Visual Studio .NET design-time support, including the Properties window and Design view previews. Also, the only way to share the user control between applications is to put a separate copy in each application, which takes more maintenance if you make changes to the control.

Web custom controls are compiled code, which makes them easier to use but more difficult to create; Web custom controls must be authored in code. Once you have created the control, however, you can add it to the Toolbox and display it in a visual designer with full Properties window support and all the other design-time features of ASP.NET server controls. In addition, you can install a single copy of the Web custom control in the global assembly cache and share it between applications, which makes maintenance easier.

Web user controls

Easier to create
Limited support for consumers who use a visual design tool
A separate copy of the control is required in each application
Cannot be added to the Toolbox in Visual Studio
Good for static layout

Web custom controls

Harder to create
Full visual design tool support for consumers
Only a single copy of the control is required, in the global assembly cache
Can be added to the Toolbox in Visual Studio
Good for dynamic layout

(Session/State)

Application and Session Events

The ASP.NET page framework provides ways for you to work with events that can be raised when your application starts or stops or when an individual user's session starts or stops:

- Application events are raised for all requests to an application. For example, **Application_BeginRequest** is raised when any Web Forms page or XML Web service in your application is requested. This event allows you to initialize resources that will be used for each request to the application. A corresponding

event, **Application_EndRequest**, provides you with an opportunity to close or otherwise dispose of resources used for the request.

- Session events are similar to application events (there is a **Session_OnStart** and a **Session_OnEnd** event), but are raised with each unique session within the application. A session begins when a user requests a page for the first time from your application and ends either when your application explicitly closes the session or when the session times out.

You can create handlers for these types of events in the Global.asax file.

Difference between ASP Session and ASP.NET Session?

asp.net session supports cookie less session & it can span across multiple servers.

What is cookie less session? How it works?

By default, ASP.NET will store the session state in the same process that processes the request, just as ASP does. If cookies are not available, a session can be tracked by adding a session identifier to the URL. This can be enabled by setting the following:

```
<sessionState cookieless="true" />
```

<http://samples.gotdotnet.com/quickstart/aspplus/doc/stateoverview.aspx>

How you will handle session when deploying application in more than a server? Describe session handling in a webfarm, how does it work and what are the limits?

By default, ASP.NET will store the session state in the same process that processes the request, just as ASP does. Additionally, ASP.NET can store session data in an external process, which can even reside on another machine. To enable this feature:

- Start the ASP.NET state service, either using the Services snap-in or by executing "net start aspnet_state" on the command line. The state service will by default listen on port 42424. To change the port, modify the registry key for the service:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\aspnet_state\Parameters\Port
- Set the **mode** attribute of the **<sessionState>** section to "StateServer".
- Configure the **stateConnectionString** attribute with the values of the machine on which you started aspnet_state.

The following sample assumes that the state service is running on the same machine as the Web server ("localhost") and uses the default port (42424):

```
<sessionState mode="StateServer" stateConnectionString="tcpip=localhost:42424" />
```

Note that if you try the sample above with this setting, you can reset the Web server (enter iisreset on the command line) and the session state value will persist.

**

What method do you use to explicitly kill a users session?

Abandon()

What are the different ways you would consider sending data across pages in ASP (i.e between 1.asp to 2.asp)?

Session

public properties

What is State Management in .Net and how many ways are there to maintain a state in .Net? What is view state?

Web pages are recreated each time the page is posted to the server. In traditional Web programming, this would ordinarily mean that all information associated with the page and the controls on the page would be lost with each round trip.

To overcome this inherent limitation of traditional Web programming, the ASP.NET page framework includes various options to help you preserve changes — that is, for managing state. The page framework includes a facility called view state that automatically preserves property values of the page and all the controls on it between round trips.

However, you will probably also have application-specific values that you want to preserve. To do so, you can use one of the state management options.

Client-Based State Management Options:

View State

Hidden Form Fields

Cookies

Query Strings

Server-Based State Management Options

Application State

Session State

Database Support

What are the disadvantages of view state / what are the benefits?

Automatic view-state management is a feature of server controls that enables them to repopulate their property values on a round trip (without you having to write any code). This feature does impact performance, however, since a server control's view state is passed to and from the server in a hidden form field. You should be aware of when view state helps you and when it hinders your page's performance.

When maintaining session through Sql server, what is the impact of Read and Write operation on Session objects? will performance degrade?

Maintaining state using database technology is a common practice when storing user-specific information where the information store is large. Database storage is particularly useful for maintaining long-term state or state that must be preserved even if the server must be restarted.

**

What are the contents of cookie?

**

How do you create a permanent cookie?

**

What is ViewState? What does the "EnableViewState" property do? Why would I want it on or off?

**

Explain the differences between Server-side and Client-side code?

Server side code will process at server side & it will send the result to client. Client side code (javascript) will execute only at client side.

Can you give an example of what might be best suited to place in the Application_Start and Session_Start subroutines?

**

Which ASP.NET configuration options are supported in the ASP.NET implementation on the shared web hosting platform?

A: Many of the ASP.NET configuration options are not configurable at the site, application or subdirectory level on the shared hosting platform.

Certain options can affect the security, performance and stability of the server and, therefore cannot be changed. The following settings are the only ones that can be changed in your site's web.config file (s):

browserCaps

clientTarget

pages

customErrors

globalization

authorization

authentication

webControls

webServices

<http://msdn.microsoft.com/library/default.asp?url=/library/en-us/cpguide/html/cpconaspnetconfiguration.asp>

Briefly describe the role of global.asax?

- How can u debug your .net application?**
- How do u deploy your asp.net application?**
- Where do we store our connection string in asp.net application?**
- Various steps taken to optimize a web based application (caching, stored procedure etc.)**
- How does ASP.NET framework maps client side events to Server side events.**

(Security)

- Security types in ASP/ASP.NET? Different Authentication modes?**
- How .Net has implemented security for web applications?**
- How to do Forms authentication in asp.net?**
- Explain authentication levels in .net ?**
- Explain auzerization levels in .net ?**
- What is Role-Based security?**

A role is a named set of principals that have the same privileges with respect to security (such as a teller or a manager). A principal can be a member of one or more roles. Therefore, applications can use role membership to determine whether a principal is authorized to perform a requested action.

**

How will you do windows authentication and what is the namespace? If a user is logged under integrated windows authentication mode, but he is still not able to logon, what might be the possible cause for this? In ASP.Net application how do you find the name of the logged in person under windows authentication?

271.What are the different authentication modes in the .NET environment?

```

272. <authentication mode="Windows|Forms|Passport|None">
273.   <forms name="name"
274.       loginUrl="url"
275.       protection="All|None|Encryption|Validation"
276.       timeout="30" path="/" >
277.       requireSSL="true|false"
278.       slidingExpiration="true|false">
279.   <credentials passwordFormat="Clear|SHA1|MD5">
280.     <user name="username" password="password"/>
281.   </credentials>
282.   </forms>
283. <passport redirectUrl="internal"/>
</authentication>

```

Attribute	Option	Description
mode		Controls the default authentication mode for an application.
	Windows	Specifies Windows authentication as the default authentication mode. Use this mode when using any form of Microsoft Internet Information Services (IIS) authentication: Basic, Digest, Integrated Windows authentication (NTLM/Kerberos), or certificates.
	Forms	Specifies ASP.NET forms-based authentication as the default authentication mode.
	Passport	Specifies Microsoft Passport authentication as the default authentication mode.
	None	Specifies no authentication. Only anonymous users are expected or applications can handle events to provide their own authentication.

How do you specify whether your data should be passed as Query string and Forms (Mainly about POST and GET)

Through attribute tag of form tag.

What is the other method, other than GET and POST, in ASP.NET?

What are validator? Name the Validation controls in asp.net? How do u disable them? Will the asp.net validators run in server side or client side? How do you do Client-side validation in .Net? How to disable validator control by client side JavaScript?

A set of server controls included with ASP.NET that test user input in HTML and Web server controls for programmer-defined requirements. Validation controls perform input checking in server code. If the user is working with a browser that supports DHTML, the validation controls can also perform validation ("EnableClientScript" property set to true/false) using client script.

The following validation controls are available in asp.net:

RequiredFieldValidator Control, CompareValidator Control, RangeValidator Control, RegularExpressionValidator Control, CustomValidator Control, ValidationSummary Control.

Which two properties are there on every validation control?

ControlToValidate, ErrorMessage

How do you use css in asp.net?

Within the <HEAD> section of an HTML document that will use these styles, add a link to this external CSS style sheet that

follows this form:

```
<LINK REL="STYLESHEET" TYPE="text/css" HREF="MyStyles.css">
```

MyStyles.css is the name of your external CSS style sheet.

How do you implement postback with a text box? What is postback and usestate?

Make AutoPostBack property to true

How can you debug an ASP page, without touching the code?

What is SQL injection?

An SQL injection attack "injects" or manipulates SQL code by adding unexpected SQL to a query.

Many web pages take parameters from web user, and make SQL query to the database. Take for instance when a user login, web page that user name and password and make SQL query to the database to check if a user has valid name and password.

Username: ' or 1=1 ---

Password: [Empty]

This would execute the following query against the users table:

```
select count(*) from users where userName="" or 1=1 --' and userPass=""
```

How can u handle Exceptions in Asp.Net?

How can u handle Un Managed Code Exceptions in ASP.Net?

Asp.net - How to find last error which occurred?

A: **Server.GetLastError();**

```
[C#]
```

```
Exception LastError;
```

```
String ErrMessage;
```

```
LastError = Server.GetLastError();
```

```
if (LastError != null)
```

```
ErrMessage = LastError.Message;
```

```
else
```

```
ErrMessage = "No Errors";
```

```
Response.Write("Last Error = " + ErrMessage);
```

How to do Caching in ASP?

A: <%@ OutputCache Duration="60" VaryByParam="None" %>

VaryByParam value	Description
none	One version of page cached (only raw GET)
*	n versions of page cached based on query string and/or

	POST body
v1	n versions of page cached based on value of v1 variable in query string or POST body
v1;v2	n versions of page cached based on value of v1 and v2 variables in query string or POST body

<%@ OutputCache Duration="60" VaryByParam="none" %>

<%@ OutputCache Duration="60" VaryByParam="*" %>

<%@ OutputCache Duration="60" VaryByParam="name;age" %>

The **OutputCache** directive supports several other cache varying options

- **VaryByHeader** - maintain separate cache entry for header string changes (**UserAgent, UserLanguage**, etc.)
- **VaryByControl** - for user controls, maintain separate cache entry for properties of a user control
- **VaryByCustom** - can specify separate cache entries for browser types and version *or* provide a custom **GetVaryByCustomString** method in **HttpApplication** derived class

297. What is the Global ASA(X) File?

Any alternative to avoid name collisions other than Namespaces.

A scenario that two namespaces named N1 and N2 are there both having the same class say A. now in another class i ve written

using N1;using N2;

and i am instantiating class A in this class. Then how will u avoid name collisions?

Ans: using alias

Eg: using MyAlias = MyCompany.Proj.Nested;

Which is the namespace used to write error message in event Log File?

What are the page level transaction and class level transaction?

What are different transaction options?

What is the namespace for encryption?

What is the difference between application and cache variables?

What is the difference between control and component?

You ve defined one page_load event in aspx page and same page_load event in code behind how will prog run?

Where would you use an IHttpModule, and what are the limitations of any approach you might take in implementing one?

307. Can you edit data in the Repeater control? Which template must you provide, in order to display data in a Repeater control? How can you provide an alternating color scheme in a Repeater control? What property must you set, and what method must you call in your code, in order to bind the data from some data source to the Repeater control?

What is the use of web.config? Difference between machine.config and Web.config?

ASP.NET configuration files are XML-based text files--each named web.config--that can appear in any directory on an ASP.NET

Web application server. Each web.config file applies configuration settings to the directory it is located in and to all

virtual child directories beneath it. Settings in child directories can optionally override or modify settings specified in

parent directories. The root configuration file--

WinNT\Microsoft.NET\Framework\<version>\config\machine.config--provides

default configuration settings for the entire machine. ASP.NET configures IIS to prevent direct browser access to web.config

files to ensure that their values cannot become public (attempts to access them will cause ASP.NET to return 403: Access

Forbidden).

At run time ASP.NET uses these web.config configuration files to hierarchically compute a unique collection of settings for

each incoming URL target request (these settings are calculated only once and then cached across subsequent requests; ASP.NET

automatically watches for file changes and will invalidate the cache if any of the

configuration files change).

<http://samples.gotdotnet.com/quickstart/aspplus/doc/configformat.aspx>

What is the use of sessionstate tag in the web.config file?

Configuring session state: Session state features can be configured via the **<sessionState>** section in a web.config file. To double the default timeout of 20 minutes, you can add the following to the web.config file of an application:

```
<sessionState
timeout="40"
/>
```

What are the different modes for the sessionstates in the web.config file?

Off	Indicates that session state is not enabled.
Inproc	Indicates that session state is stored locally.
StateServer	Indicates that session state is stored on a remote server.
SQLServer	Indicates that session state is stored on the SQL Server.

What is smart navigation?

When a page is requested by an Internet Explorer 5 browser, or later, smart navigation enhances the user's experience of the page by performing the following:

- eliminating the flash caused by navigation.
- persisting the scroll position when moving from page to page.
- persisting element focus between navigations.
- retaining only the last page state in the browser's history.

Smart navigation is best used with ASP.NET pages that require frequent postbacks but with visual content that does not change dramatically on return. Consider this carefully when deciding whether to set this property to **true**.

Set the **SmartNavigation** attribute to **true** in the @ Page directive in the .aspx file.

When the page is requested, the dynamically generated class sets this property.

In what order do the events of an ASPX page execute. As a developer is it important to understand these events?

313. How would you get ASP.NET running in Apache web servers - why would you even do this?

What tags do you need to add within the asp:datagrid tags to bind columns manually

What base class do all Web Forms inherit from?

System.Web.UI.Page

How can we create pie chart in asp.net?

Is it possible for me to change my aspx file extension to some other name?

Yes.

Open IIS->Default Website -> Properties

Select HomeDirectory tab

Click on configuration button

Click on add. Enter aspnet_isapi details

(C:\WINDOWS\Microsoft.NET\Framework\v1.0.3705\aspnet_isapi.dll | GET,HEAD,POST,DEBUG)

Open machine.config(C:\WINDOWS\Microsoft.NET\Framework\v1.0.3705\CONFIG) & add new extension under <httpHandlers> tag

```
<add verb="*" path="*.santhosh" type="System.Web.UI.PageHandlerFactory"/>
```

What is AutoEventWireup attribute for ?

(WEBSERVICE & REMOTING)

What is a WebService and what is the underlying protocol used in it? Why Web Services?

Web Services are applications delivered as a service on the Web. Web services allow for programmatic access of business logic over the Web. Web services typically rely on XML-based protocols, messages, and interface descriptions for communication and access. Web services are designed to be used by other programs or applications rather

than directly by end user. Programs invoking a Web service are called clients. SOAP over HTTP is the most commonly used protocol for invoking Web services. There are three main uses of Web services.

Application integration Web services within an intranet are commonly used to integrate business applications running on disparate platforms. For example, a .NET client running on Windows 2000 can easily invoke a Java Web service running on a mainframe or Unix machine to retrieve data from a legacy application.

Business integration Web services allow trading partners to engage in e-business leveraging the existing Internet infrastructure. Organizations can send electronic purchase orders to suppliers and receive electronic invoices. Doing e-business with Web services means a low barrier to entry because Web services can be added to existing applications running on any platform without changing legacy code.

3. Commercial Web services focus on selling content and business services to clients over the Internet similar to familiar Web pages. Unlike Web pages, commercial Web services target applications not humans as their direct users. Continental Airlines exposes flight schedules and status Web services for travel Web sites and agencies to use in their applications. Like Web pages, commercial Web services are valuable only if they expose a valuable service or content. It would be very difficult to get customers to pay you for using a Web service that creates business charts with the customers' data. Customers would rather buy a charting component (e.g. COM or .NET component) and install it on the same machine as their application. On the other hand, it makes sense to sell real-time weather information or stock quotes as a Web service. Technology can help you add value to your services and explore new markets, but ultimately customers pay for contents and/or business services, not for technology

Are Web Services a replacement for other distributed computing platforms?

No. Web Services is just a new way of looking at existing implementation platforms.

In a Webservice, need to display 10 rows from a table. So DataReader or DataSet is best choice?

A: WebService will support only DataSet.

How to generate WebService proxy? What is SOAP, WSDL, UDDI and the concept behind Web Services? What are various components of WSDL? What is the use of WSDL.exe utility?

SOAP is an XML-based messaging framework specifically designed for exchanging formatted data across the Internet, for example using request and reply messages or sending entire documents. SOAP is simple, easy to use, and completely neutral with respect to operating system, programming language, or distributed computing platform.

After SOAP became available as a mechanism for exchanging XML messages among enterprises (or among disparate applications within the same enterprise), a better way was needed to describe the messages and how they are exchanged. The Web Services Description Language (WSDL) is a particular form of an XML Schema, developed by Microsoft and IBM for the purpose of defining the XML message, operation, and protocol mapping of a web service accessed using SOAP or other XML protocol. WSDL defines web services in terms of "endpoints" that operate on XML messages. The WSDL syntax allows both the messages and the operations on the messages to be defined abstractly, so they can be mapped to multiple physical implementations. The current WSDL spec describes how to map messages and operations to SOAP 1.1, HTTP GET/POST, and MIME. WSDL creates web service definitions by mapping a group of endpoints into a logical sequence of operations on XML messages. The same XML message can be mapped to multiple operations (or services) and bound to one or more communications protocols (using "ports").

The Universal Description, Discovery, and Integration (UDDI) framework defines a data model (in XML) and SOAP APIs for registration and searches on business information, including the web services a business exposes to the Internet. UDDI is an independent consortium of vendors, founded by Microsoft, IBM, and Ariba, for the purpose of

developing an Internet standard for web service description registration and discovery. Microsoft, IBM, and Ariba also are hosting the initial deployment of a UDDI service, which is conceptually patterned after DNS (the Internet service that translates URLs into TCP addresses). UDDI uses a private agreement profile of SOAP (i.e. UDDI doesn't use the SOAP serialization format because it's not well suited to passing complete XML documents (it's aimed at RPC style interactions). The main idea is that businesses use the SOAP APIs to register themselves with UDDI, and other businesses search UDDI when they want to discover a trading partner, for example someone from whom they wish to procure sheet metal, bolts, or transistors. The information in UDDI is categorized according to industry type and geographical location, allowing UDDI consumers to search through lists of potentially matching businesses to find the specific one they want to contact. Once a specific business is chosen, another call to UDDI is made to obtain the specific contact information for that business. The contact information includes a pointer to the target business's WSDL or other XML schema file describing the web service that the target business publishes.

How to generate proxy class other than .net app and wsdl tool?

To access an XML Web service from a client application, you first add a Web reference, which is a reference to an XML Web service. When you create a Web reference, Visual Studio creates an XML Web service proxy class automatically and adds it to your project. This proxy class exposes the methods of the XML Web service and handles the marshalling of appropriate arguments back and forth between the XML Web service and your application. Visual Studio uses the Web Services Description Language (WSDL) to create the proxy.

To generate an XML Web service proxy class:

- From a command prompt, use Wsdl.exe to create a proxy class, specifying (at a minimum) the URL to an XML Web service or a service description, or the path to a saved service description.

```
Wsdl /language:language /protocol:protocol /namespace:myNameSpace  
/out:filename  
/username:username /password:password /domain:domain <url or path>
```

What is a proxy in web service? How do I use a proxy server when invoking a Web service?

asynchronous web service means?

What are the events fired when web service called?

How will do transaction in Web Services?

How does SOAP transport happen and what is the role of HTTP in it? How you can access a webservice using soap?

What are the different formatters can be used in both? Why?.. binary/soap

How you will protect / secure a web service?

For the most part, things that you do to secure a Web site can be used to secure a Web Service. If you need to encrypt the data exchange, you use Secure Sockets Layer (SSL) or a Virtual Private Network to keep the bits secure. For authentication, use HTTP Basic or Digest authentication with Microsoft® Windows® integration to figure out who the caller is.

these items cannot:

- Parse a SOAP request for valid values
- Authenticate access at the Web Method level (they can authenticate at the Web Service level)
- Stop reading a request as soon as it is recognized as invalid

<http://msdn.microsoft.com/library/default.asp?url=/library/en-us/cpguide/html/cpcontransactionsupportinaspnetwebservices.asp>

How will you expose/publish a webservice?

What is disco file?

What's the attribute for webservice method? What is the namespace for creating webservice?

[WebMethod]

using System.Web;
using System.Web.Services;

What is Remoting?

The process of communication between different operating system processes, regardless of whether they are on the same computer. The .NET remoting system is an architecture designed to simplify communication between objects living in different application domains, whether on the same computer or not, and between different contexts, whether in the same application domain or not.

Difference between web services & remoting?

	ASP.NET Web Services	.NET Remoting
Protocol	Can be accessed only over HTTP	Can be accessed over any protocol (including TCP, HTTP, SMTP and so on)
State Management	Web services work in a stateless environment	Provide support for both stateful and stateless environments through Singleton and SingleCall objects
Type System	Web services support only the datatypes defined in the XSD type system, limiting the number of objects that can be serialized.	Using binary communication, .NET Remoting can provide support for rich type system
Interoperability	Web services support interoperability across platforms, and are ideal for heterogeneous environments.	.NET remoting requires the client be built using .NET, enforcing homogenous environment.
Reliability	Highly reliable due to the fact that Web services are always hosted in IIS	Can also take advantage of IIS for fault isolation. If IIS is not used, application needs to provide plumbing for ensuring the reliability of the application.
Extensibility	Provides extensibility by allowing us to intercept the SOAP messages during the serialization and deserialization stages.	Very extensible by allowing us to customize the different components of the .NET remoting framework.
Ease-of-Programming	Easy-to-create and deploy.	Complex to program.

Though both the .NET Remoting infrastructure and ASP.NET Web services can enable cross-process communication, each is designed to benefit a different target audience. ASP.NET Web services provide a simple programming model and a wide reach. .NET Remoting provides a more complex programming model and has a much narrower reach.

As explained before, the clear performance advantage provided by TCPChannel-remoting should make you think about using this channel whenever you can afford to do so. If you can create direct TCP connections from your clients to your server and if you need to support only the .NET platform, you should go for this channel. If you are going to go cross-platform or you have the requirement of supporting SOAP via HTTP, you should definitely go for ASP.NET Web services.

Both the .NET remoting and ASP.NET Web services are powerful technologies that provide a suitable framework for developing distributed applications. It is important to understand how both technologies work and then choose the one that is right for your application. For applications that require interoperability and must function over public networks, Web services are probably the best bet. For those that require

communications with other .NET components and where performance is a key priority, .NET Remoting is the best choice. In short, use Web services when you need to send and receive data from different computing platforms, use .NET Remoting when sending and receiving data between .NET applications. In some architectural scenarios, you might also be able to use .NET Remoting in conjunction with ASP.NET Web services and take advantage of the best of both worlds.

The Key difference between ASP.NET webservice and .NET Remoting is how they serialize data into messages and the format they choose for metadata. ASP.NET uses XML serializer for serializing or Marshalling. And XSD is used for Metadata. .NET Remoting relies on System.Runtime.Serialization.Formatter.Binary and System.Runtime.Serialization.SOAPFormatter and relies on .NET CLR Runtime assemblies for metadata.

Can you pass SOAP messages through remoting?

CAO and SAO.

Client Activated objects are those remote objects whose Lifetime is directly Controlled by the client. This is in direct contrast to SAO. Where the server, not the client has complete control over the lifetime of the objects.

Client activated objects are instantiated on the server as soon as the client request the object to be created. Unlike as SAO a CAO doesn't delay the object creation until the first method is called on the object. (In SAO the object is instantiated when the client calls the method on the object)

singleton and singlecall.

Singleton types never have more than one instance at any one time. If an instance exists, all client requests are serviced by that instance.

Single Call types always have one instance per client request. The next method invocation will be serviced by a different server instance, even if the previous instance has not yet been recycled by the system.

18. What is Asynchronous Web Services?

Web Client class and its methods?

Flow of remoting?

What is the use of trace utility?

Using the SOAP Trace Utility

The Microsoft® Simple Object Access Protocol (SOAP) Toolkit 2.0 includes a TCP/IP trace utility, MSSOAPT.EXE. You use this trace utility to view the SOAP messages sent by HTTP between a SOAP client and a service on the server.

Using the Trace Utility on the Server

To see all of a service's messages received from and sent to all clients, perform the following steps on the server.

1. On the server, open the Web Services Description Language (WSDL) file. In the WSDL file, locate the <soap:address> element that corresponds to the service and change the **location** attribute for this element to port 8080. For example, if the **location** attribute specifies <http://MyServer/VDir/Service.wsdl> change this attribute to <http://MyServer:8080/VDir/Service.wsdl>. Run **MSSOAPT.exe**.
On the **File** menu, point to **New**, and either click **Formatted Trace** (if you don't want to see HTTP headers) or click **Unformatted Trace** (if you do want to see HTTP headers).
In the **Trace Setup** dialog box, click **OK** to accept the default values.

Using the Trace Utility on the Client

To see all messages sent to and received from a service, do the following steps on the client.

6. Copy the WSDL file from the server to the client.
7. Modify location attribute of the <soap:address> element in the local copy of the WSDL document to direct the client to localhost:8080 and make a note of the

current host and port. For example, if the WSDL contains
<http://MyServer/VDir/Service.wsdl>, change it to
<http://localhost:8080/VDir/Service.wsdl> and make note of "MyServer".
On the client, run **MSSOPT.exe**.

On the **File** menu, point to **New**, and either click **Formatted Trace** (if you don't want to see HTTP headers) or click **Unformatted Trace** (if you do want to see HTTP headers).

In the **Destination host** box, enter the host specified in Step 2.

In the **Destination port** box, enter the port specified in Step 2.

Click **OK**.

(XML)

Explain the concept of data island?

How to use XML DOM model on client side using JavaScript.

What are the ways to create a tree view control using XML, XSL & JavaScript?

Questions on XPathNavigator, and the other classes in System.XML

Namespace?

What is Use of Template in XSL?

What is "Well Formed XML" and "Valid XML"

How you will do SubString in XSL

Can we do sorting in XSL ? how do you deal sorting columns dynamically in XML.

What is "Async" property of XML Means ?

What is XPath Query ?

Difference Between Element and Node.

What is CDATA Section.

DOM & SAX parsers explanation and difference

What is GetElementbyname method will do?

What is selectnode method will give?

85. What is valid xml document? What a well formed xml document?

86. What is the Difference between XmlDocument and XmlDataDocument?

Explain what a DiffGram is, and a good use for one?

A DiffGram is an XML format that is used to identify current and original versions of data elements. When sending and retrieving a **DataSet** from an XML Web service, the DiffGram format is implicitly used.

The **DataSet** uses the DiffGram format to load and persist its contents, and to serialize its contents for transport across a network connection. When a **DataSet** is written as a DiffGram, it populates the DiffGram with all the necessary information to accurately recreate the contents, though not the schema, of the **DataSet**, including column values from both the **Original** and **Current** row versions, row error information, and row order.

DiffGram Format

The DiffGram format is divided into three sections: the current data, the original (or "before") data, and an errors section, as shown in the following example.

```
<?xml version="1.0"?>
<diffgr:diffgram
xmlns:msdata="urn:schemas-microsoft-com:xml-msdata"
xmlns:diffgr="urn:schemas-microsoft-com:xml-diffgram-v1"
xmlns:xsd="http://www.w3.org/2001/XMLSchema">
```

```
<DataInstance>
</DataInstance>
```

```
<diffgr:before>
</diffgr:before>
```

```
<diffgr:errors>
</diffgr:errors>
</diffgr:diffgram>
```

The DiffGram format consists of the following blocks of data:

<DataInstance>

The name of this element, **DataInstance**, is used for explanation purposes in this documentation. A **DataInstance** element represents a **DataSet** or a row of a **DataTable**. Instead of *DataInstance*, the element would contain the name of the **DataSet** or **DataTable**. This block of the DiffGram format contains the current data, whether it has been modified or not. An element, or row, that has been modified is identified with the **diffgr:hasChanges** annotation.

<diffgr:before>

This block of the DiffGram format contains the original version of a row. Elements in this block are matched to elements in the **DataInstance** block using the **diffgr:id** annotation.

<diffgr:errors>

This block of the DiffGram format contains error information for a particular row in the **DataInstance** block. Elements in this block are matched to elements in the **DataInstance** block using the **diffgr:id** annotation.

**If I replace my Sqlserver with XML files and how about handling the same?
Write syntax to serialize class using XML Serializer?**

(IIS)

In which process does IIS runs (was asking about the EXE file)

inetinfo.exe is the Microsoft IIS server running, handling ASP.NET requests among other things. When an ASP.NET request is received (usually a file with .aspx extension), the ISAPI filter aspnet_isapi.dll takes care of it by passing the request to the actual worker process aspnet_wp.exe.

Where are the IIS log files stored?

C:\WINDOWS\system32\Logfiles\W3SVC1

OR

c:\winnt\system32\LogFiles\W3SVC1

What are the different IIS authentication modes in IIS 5.0 and Explain?

Difference between basic and digest authentication modes?

IIS provides a variety of authentication schemes:

- Anonymous (enabled by default)
- Basic
- Digest
- Integrated Windows authentication (enabled by default)
- Client Certificate Mapping

Anonymous

Anonymous authentication gives users access to the public areas of your Web site without prompting them for a user name or password. Although listed as an authentication scheme, it is not technically performing any client authentication because the client is not required to supply any credentials. Instead, IIS provides stored credentials to Windows using a special user account, IUSR_ *machinename*. By default, IIS controls the password for this account. Whether or not IIS controls the password affects the permissions the anonymous user has. When IIS controls the password, a sub authentication DLL (iissuba.dll) authenticates the user using a network logon. The function of this DLL is to validate the password supplied by IIS and to inform Windows that the password is valid, thereby authenticating the client. However, it does not actually provide a password to Windows. When IIS does not control the password, IIS calls the LogonUser() API in Windows and provides the account name, password

and domain name to log on the user using a local logon. After the logon, IIS caches the security token and impersonates the account. A local logon makes it possible for the anonymous user to access network resources, whereas a network logon does not.

Basic Authentication

IIS Basic authentication as an implementation of the basic authentication scheme found in section 11 of the [HTTP 1.0 specification](#).

As the specification makes clear, this method is, in and of itself, non-secure. The reason is that Basic authentication assumes a trusted connection between client and server.

Thus, the username and password are transmitted in clear text. More specifically, they are transmitted using Base64 encoding, which is trivially easy to decode. This makes Basic authentication the wrong choice to use over a public network on its own.

Basic Authentication is a long-standing standard supported by nearly all browsers. It also imposes no special requirements on the server side -- users can authenticate against any NT domain, or even against accounts on the local machine. With SSL to shelter the security credentials while they are in transmission, you have an authentication solution that is both highly secure and quite flexible.

Digest Authentication

The Digest authentication option was added in Windows 2000 and IIS 5.0. Like Basic authentication, this is an implementation of a technique suggested by Web standards, namely [RFC 2069](#) (superceded by [RFC 2617](#)).

Digest authentication also uses a challenge/response model, but it is much more secure than Basic authentication (when used without SSL). It achieves this greater security not by encrypting the secret (the password) before sending it, but rather by following a different design pattern -- one that does not require the client to transmit the password over the wire at all.

Instead of sending the password itself, the client transmits a one-way message digest (a checksum) of the user's password, using (by default) the MD5 algorithm. The server then fetches the password for that user from a Windows 2000 Domain Controller, reruns the checksum algorithm on it, and compares the two digests. If they match, the server knows that the client knows the correct password, even though the password itself was never sent. (If you have ever wondered what the default ISAPI filter "md5filt" that is installed with IIS 5.0 is used for, now you know.

Integrated Windows Authentication

Integrated Windows authentication (formerly known as NTLM authentication and Windows NT Challenge/Response authentication) can use either NTLM or Kerberos V5 authentication and only works with Internet Explorer 2.0 and later.

When Internet Explorer attempts to access a protected resource, IIS sends two WWW-Authenticate headers, Negotiate and NTLM.

- If Internet Explorer recognizes the Negotiate header, it will choose it because it is listed first. When using Negotiate, the browser will return information for both NTLM and Kerberos. At the server, IIS will use Kerberos if both the client (Internet Explorer 5.0 and later) and server (IIS 5.0 and later) are running Windows 2000 and later, and both are members of the same domain or trusted domains. Otherwise, the server will default to using NTLM.
- If Internet Explorer does not understand Negotiate, it will use NTLM.

So, which mechanism is used depends upon a negotiation between Internet Explorer and IIS.

When used in conjunction with Kerberos v5 authentication, IIS can delegate security credentials among computers running Windows 2000 and later that are trusted and configured for delegation. Delegation enables remote access of resources on behalf of the delegated user.

Integrated Windows authentication is the best authentication scheme in an intranet environment where users have Windows domain accounts, especially when using Kerberos. Integrated Windows authentication, like digest authentication, does not pass the user's password across the network. Instead, a hashed value is exchanged.

Client Certificate Mapping

A certificate is a digitally signed statement that contains information about an entity

and the entity's public key, thus binding these two pieces of information together. A trusted organization (or entity) called a Certification Authority (CA) issues a certificate after the CA verifies that the entity is who it says it is. Certificates can contain different types of data. For example, an X.509 certificate includes the format of the certificate, the serial number of the certificate, the algorithm used to sign the certificate, the name of the CA that issued the certificate, the name and public key of the entity requesting the certificate, and the CA's signature. X.509 client certificates simplify authentication for larger user bases because they do not rely on a centralized account database. You can verify a certificate simply by examining the certificate.
<http://msdn.microsoft.com/library/default.asp?url=/library/en-us/vsent7/html/vxconIISAuthentication.asp>

How to configure the sites in Web server (IIS)?

Advantages in IIS 6.0?

<http://www.microsoft.com/windowsserver2003/iis/evaluation/features/default.msp>

http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windowsserver2003/proddocs/datacenter/gs_whatschanged.asp

IIS Isolation Levels?

Internet Information Server introduced the notion "Isolation Level", which is also present in IIS4 under a different name. IIS5 supports three isolation levels, that you can set from the Home Directory tab of the site's Properties dialog:

- **Low (IIS Process):** ASP pages run in INetInfo.Exe, the main IIS process, therefore they are executed in-process. This is the fastest setting, and is the default under IIS4. The problem is that if ASP crashes, IIS crashes as well and must be restarted (IIS5 has a reliable restart feature that automatically restarts a server when a fatal error occurs).
- **Medium (Pooled):** In this case ASP runs in a different process, which makes this setting more reliable: if ASP crashes IIS won't. All the ASP applications at the Medium isolation level share the same process, so you can have a web site running with just two processes (IIS and ASP process). IIS5 is the first Internet Information Server version that supports this setting, which is also the default setting when you create an IIS5 application. Note that an ASP application that runs at this level is run under COM+, so it's hosted in DLLHOST.EXE (and you can see this executable in the Task Manager).
- **High (Isolated):** Each ASP application runs out-process in its own process space, therefore if an ASP application crashes, neither IIS nor any other ASP application will be affected. The downside is that you consume more memory and resources if the server hosts many ASP applications. Both IIS4 and IIS5 supports this setting: under IIS4 this process runs inside MTS.EXE, while under IIS5 it runs inside DLLHOST.EXE.

When selecting an isolation level for your ASP application, keep in mind that out-process settings - that is, Medium and High - are less efficient than in-process (Low). However, out-process communication has been vastly improved under IIS5, and in fact IIS5's Medium isolation level often deliver better results than IIS4's Low isolation. In practice, you shouldn't set the Low isolation level for an IIS5 application unless you really need to serve hundreds pages per second.

Controls

How will you do Redo and Undo in a TextControl?

How to implement DataGrid in .NET? How would u make a combo-box appear in one column of a DataGrid? What are the ways to show data grid inside a data grid for a master details type of tables? If we write any code for DataGrid methods, what is the access specifier used for that methods in the code behind file and why?

How can we create Tree control in asp.net?

Programming

Write a program in C# for checking a given number is PRIME or not.

Write a program to find the angle between the hours and minutes in a clock

Write a C# program to find the Factorial of n

How do I upload a file from my ASP.NET page?

A: In order to perform file upload in your ASP.NET page, you will need to use two classes: the System.Web.UI.HtmlControls.HtmlInputFile class and the System.Web.HttpPostedFile class. The HtmlInputFile class represents an HTML input control that the user will use on the client side to select a file to upload. The HttpPostedFile class represents the uploaded file and is obtained from the PostedFile property of the HtmlInputFile class. In order to use the HtmlInputFile control, you need to add the enctype attribute to your form tag as follows:

```
<form id="upload" method="post" runat="server" enctype="multipart/form-data">
```

Also, remember that the /data directory is the only directory with Write permissions enabled for the anonymous user. Therefore, you will need to make sure that your code uploads the file to the /data directory or one of its subdirectories.

Below is a simple example of how to upload a file via an ASP.NET page in C# and VB.NET.

C#

```
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Web" %>
<%@ Import Namespace="System.Web.UI.HtmlControls" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Drawing" %>
<html>
<head>
<title>upload_cs</title>
</head>
<script language="C#" runat="server">
public void UploadFile(object sender, EventArgs e)
{
if (loFile.PostedFile != null)
{
try
{
string strFileName, strFileNamePath, strFileFolder;
strFileFolder = Context.Server.MapPath(@"data\");
strFileName = loFile.PostedFile.FileName;
strFileName = Path.GetFileName(strFileName);
strFileNamePath = strFileFolder + strFileName;
loFile.PostedFile.SaveAs(strFileNamePath);
lblFileName.Text = strFileName;
lblFileLength.Text = loFile.PostedFile.ContentLength.ToString();
lblFileType.Text = loFile.PostedFile.ContentType;
pnStatus.Visible = true;
}
catch (Exception x)
{
Label lblError = new Label();
lblError.ForeColor = Color.Red;
lblError.Text = "Exception occurred: " + x.Message;
lblError.Visible = true;
this.Controls.Add(lblError);
}
}
}
}
</script>
```

```

<body>
<form id="upload_cs" method="post" runat="server" enctype="multipart/form-data">
<P>
<INPUT type="file" id="loFile" runat="server">
</P>
<P>
<asp:Button id="btnUpload" runat="server" Text=" Upload "
OnClick="UploadFile"></asp:Button></P>
<P>
<asp:Panel id="pnStatus" runat="server" Visible="False">
<asp:Label id="lblFileName" Font-Bold="True" Runat="server"></asp:Label>
uploaded<BR>
<asp:Label id="lblFileLength" Runat="server"></asp:Label> bytes<BR>
<asp:Label id="lblFileType" Runat="server"></asp:Label>
</asp:Panel></P>
</form>
</body>
</html>

```

How do I send an email message from my ASP.NET page?

A: You can use the System.Web.Mail.MailMessage and the System.Web.Mail.SmtpMail class to send email in your ASPX pages. Below is a simple example of using this class to send mail in C# and VB.NET. In order to send mail through our mail server, you would want to make sure to set the static SmtpServer property of the SmtpMail class to mail-fwd.

C#

```

<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Web" %>
<%@ Import Namespace="System.Web.Mail" %>
<HTML>
<HEAD>
<title>Mail Test</title>
</HEAD>
<script language="C#" runat="server">
private void Page_Load(Object sender, EventArgs e)
{
try
{
MailMessage mailObj = new MailMessage();
mailObj.From = "sales@joeswidgets.com";
mailObj.To = "ringleader@forexample-domain.com";
mailObj.Subject = "Your Widget Order";
mailObj.Body = "Your order was processed.";
mailObj.BodyFormat = MailFormat.Text;
SmtpMail.SmtpServer = "mail-fwd";
SmtpMail.Send(mailObj);
Response.Write("Mail sent successfully");
}
catch (Exception x)
{
Response.Write("Your message was not sent: " + x.Message);
}
}
</script>
<body>
<form id="mail_test" method="post" runat="server">
</form>
</body>
</HTML>

```

**13. Write a program to create a user control with name and surname as data members and login as method and also the code to call it. (Hint use event delegates) Practical Example of Passing an Events to delegates
How can you read 3rd line from a text file?**

Areas for study

Assemblies, GAC (how to post private assembly to gac)

.net architecture, MSIL, CTS, CLR

Events, delegates (this is the basics of .net. u have to understand it very well)

asp.net, webform, server controls, user controls

ado.net, dataset, datareader, dataadapter

remoting, webservice

desktop application - datagrid.

Threading

.NET Framework Frequently Asked Questions

Andy McMullan

Last update: 5-Aug-2004

This FAQ tries to answer some commonly asked questions about the fundamentals of the .NET Framework - topics like assemblies, garbage collection, security, interop with COM, and remoting. The most commonly-used parts of the class library are also covered. Other aspects of the .NET Framework such as ASP.NET, ADO.NET and WinForms are not covered.

This FAQ was inspired by discussions on the DOTNET mailing list. The list has now been split into several DOTNET-X lists - for details see <http://discuss.develop.com/>.

Christophe Lauer has translated the FAQ into French - you can find it at http://www.dotnet-fr.org/documents/andy_faqqdotnet_fr.html

.

Contents

- 1. Introduction
 - 1.1 What is .NET?
 - 1.2 Does .NET only apply to people building web-sites?
 - 1.3 When was .NET announced?
 - 1.4 When was the first version of .NET released?
 - 1.5 What tools can I use to develop .NET applications?
 - 1.6 What platforms does the .NET Framework run on?
 - 1.7 What languages does the .NET Framework support?

- 1.8 Will the .NET Framework go through a standardisation process?
- 2. Basic terminology
 - 2.1 What is the CLR?
 - 2.2 What is the CTS?
 - 2.3 What is the CLS?
 - 2.4 What is IL?
 - 2.5 What is C#?
 - 2.6 What does 'managed' mean in the .NET context?
 - 2.7 What is reflection?
- 3. Assemblies
 - 3.1 What is an assembly?
 - 3.2 How can I produce an assembly?
 - 3.3 What is the difference between a private assembly and a shared assembly?
 - 3.4 How do assemblies find each other?
 - 3.5 How does assembly versioning work?
- 4. Application Domains
 - 4.1 What is an Application Domain?
 - 4.2 How does an AppDomain get created?
 - 4.3 Can I write my own .NET host?
- 5. Garbage Collection
 - 5.1 What is garbage collection?
 - 5.2 Is it true that objects don't always get destroyed immediately when the last reference goes away?
 - 5.3 Why doesn't the .NET runtime offer deterministic destruction?
 - 5.4 Is the lack of deterministic destruction in .NET a problem?
 - 5.5 Does non-deterministic destruction affect the usage of COM objects from managed code?
 - 5.6 I've heard that Finalize methods should be avoided. Should I implement Finalize on my class?
 - 5.7 Do I have any control over the garbage collection algorithm?
 - 5.8 How can I find out what the garbage collector is doing?
- 6. Serialization
 - 6.1 What is serialization?
 - 6.2 Does the .NET Framework have in-built support for serialization?
 - 6.3 I want to serialize instances of my class. Should I use XmlSerializer, SoapFormatter or BinaryFormatter?
 - 6.4 Can I customise the serialization process?
 - 6.5 Why is XmlSerializer so slow?
 - 6.6 Why do I get errors when I try to serialize a Hashtable?

- 6.7 XmlSerializer is throwing a generic "There was an error reflecting MyClass" error. How do I find out what the problem is?
- 7. Attributes
 - 7.1 What are attributes?
 - 7.2 Can I create my own metadata attributes?
 - 7.3 Can I create my own context attributes?
- 8. Code Access Security
 - 8.1 What is Code Access Security (CAS)?
 - 8.2 How does CAS work?
 - 8.3 Who defines the CAS code groups?
 - 8.4 How do I define my own code group?
 - 8.5 How do I change the permission set for a code group?
 - 8.6 Can I create my own permission set?
 - 8.7 I'm having some trouble with CAS. How can I diagnose my problem?
 - 8.8 I can't be bothered with all this CAS stuff. Can I turn it off?
- 9. Intermediate Language (IL)
 - 9.1 Can I look at the IL for an assembly?
 - 9.2 Can source code be reverse-engineered from IL?
 - 9.3 How can I stop my code being reverse-engineered from IL?
 - 9.4 Can I write IL programs directly?
 - 9.5 Can I do things in IL that I can't do in C#?
- 10. Implications for COM
 - 10.1 Is COM dead?
 - 10.2 Is DCOM dead?
 - 10.3 Is MTS/COM+ dead?
 - 10.4 Can I use COM components from .NET programs?
 - 10.5 Can I use .NET components from COM programs?
 - 10.6 Is ATL redundant in the .NET world?
- 11. Miscellaneous
 - 11.1 How does .NET remoting work?
 - 11.2 How can I get at the Win32 API from a .NET program?
- 12. Class Library
 - 12.1 File I/O
 - 12.1.1 How do I read from a text file?
 - 12.1.2 How do I write to a text file?
 - 12.1.3 How do I read/write binary files?
 - 12.2 Text Processing
 - 12.2.1 Are regular expressions supported?
 - 12.3 Internet
 - 12.3.1 How do I download a web page?
 - 12.3.2 How do I use a proxy?

- 12.4 XML
 - 12.4.1 Is DOM supported?
 - 12.4.2 Is SAX supported?
 - 12.4.3 Is XPath supported?
- 12.5 Threading
 - 12.5.1 Is multi-threading supported?
 - 12.5.2 How do I spawn a thread?
 - 12.5.3 How do I stop a thread?
 - 12.5.4 How do I use the thread pool?
 - 12.5.5 How do I know when my thread pool work item has completed?
 - 12.5.6 How do I prevent concurrent access to my data?
- 12.6 Tracing
 - 12.6.1 Is there built-in support for tracing/logging?
 - 12.6.2 Can I redirect tracing to a file?
 - 12.6.3 Can I customise the trace output?
- 13. Resources
 - 13.1 Recommended books
 - 13.2 Internet Resources
 - 13.3 Weblogs
 - 13.4 Sample code & utilities

1. Introduction

1.1 What is .NET?

That's difficult to sum up in a sentence. According to Microsoft, .NET is a "revolutionary new platform, built on open Internet protocols and standards, with tools and services that meld computing and communications in new ways".

A more practical definition would be that .NET is a new environment for developing and running software applications, featuring ease of development of web-based services, rich standard run-time services available to components written in a variety of programming languages, and inter-language and inter-machine interoperability.

Note that when the term ".NET" is used in this FAQ it refers only to the new .NET runtime and associated technologies. This is sometimes called the ".NET Framework". This FAQ does NOT cover any of the various other existing and new products/technologies that Microsoft are attaching the .NET name to (e.g. SQL Server.NET).

1.2 Does .NET only apply to people building web-sites?

No. If you write any Windows software (using ATL/COM, MFC, VB, or even raw Win32), .NET may offer a viable alternative (or addition) to the way you do

things currently. Of course, if you *do* develop web sites, then .NET has lots to interest you - not least ASP.NET.

1.3 When was .NET announced?

Bill Gates delivered a keynote at Forum 2000, held June 22, 2000, outlining the .NET 'vision'. The July 2000 PDC had a number of sessions on .NET technology, and delegates were given CDs containing a pre-release version of the .NET framework/SDK and Visual Studio.NET.

1.4 When was the first version of .NET released?

The final version of the 1.0 SDK and runtime was made publicly available around 6pm PST on 15-Jan-2002. At the same time, the final version of Visual Studio.NET was made available to MSDN subscribers.

1.5 What tools can I use to develop .NET applications?

There are a number of tools, described here in ascending order of cost:

- .NET Framework SDK: The SDK is free and includes command-line compilers for C++, C#, and VB.NET and various other utilities to aid development.
- ASP.NET Web Matrix: This is a free ASP.NET development environment from Microsoft. As well as a GUI development environment, the download includes a simple web server that can be used instead of IIS to host ASP.NET apps. This opens up ASP.NET development to users of Windows XP Home Edition, which cannot run IIS.
- Microsoft Visual C# .NET Standard 2003: This is a cheap (around \$100) version of Visual Studio limited to one language and also with limited wizard support. For example, there's no wizard support for class libraries or custom UI controls. Useful for beginners to learn with, or for savvy developers who can work around the deficiencies in the supplied wizards. As well as C#, there are VB.NET and C++ versions.
- Microsoft Visual Studio.NET Professional 2003: If you have a license for Visual Studio 6.0, you can get the upgrade. You can also upgrade from VS.NET 2002 for a token \$30. Visual Studio.NET includes support for all the MS languages (C#, C++, VB.NET) and has extensive wizard support.

At the top end of the price spectrum are the Visual Studio.NET 2003 Enterprise and Enterprise Architect editions. These offer extra features such as Visual Sourcesafe (version control), and performance and analysis tools. Check out the Visual Studio.NET Feature Comparison at <http://msdn.microsoft.com/vstudio/howtobuy/choosing.asp>.

1.6 What platforms does the .NET Framework run on?

The runtime supports Windows XP, Windows 2000, NT4 SP6a and Windows ME/98. Windows 95 is not supported. Some parts of the framework do not

work on all platforms - for example, ASP.NET is only supported on Windows XP and Windows 2000. Windows 98/ME cannot be used for development.

IIS is not supported on Windows XP Home Edition, and so cannot be used to host ASP.NET. However, the [ASP.NET Web Matrix](#) web server **does** run on XP Home.

The [Mono](#) project is attempting to implement the .NET framework on Linux.

1.7 What languages does the .NET Framework support?

MS provides compilers for C#, C++, VB and JScript. Other vendors have announced that they intend to develop .NET compilers for languages such as COBOL, Eiffel, Perl, Smalltalk and Python.

1.8 Will the .NET Framework go through a standardisation process?

From <http://msdn.microsoft.com/net/ecma/>: "On December 13, 2001, the ECMA General Assembly ratified the C# and common language infrastructure (CLI) specifications into international standards. The ECMA standards will be known as ECMA-334 (C#) and ECMA-335 (the CLI)."

2. Basic terminology

2.1 What is the CLR?

CLR = Common Language Runtime. The CLR is a set of standard resources that (in theory) any .NET program can take advantage of, regardless of programming language. Robert Schmidt (Microsoft) lists the following CLR resources in his MSDN PDC# [article](#):

- Object-oriented programming model (inheritance, polymorphism, exception handling, garbage collection)
- Security model
- Type system
- All .NET base classes
- Many .NET framework classes
- Development, debugging, and profiling tools
- Execution and code management
- IL-to-native translators and optimizers

What this means is that in the .NET world, different programming languages will be more equal in capability than they have ever been before, although clearly not *all* languages will support *all* CLR services.

2.2 What is the CTS?

CTS = Common Type System. This is the range of types that the .NET runtime understands, and therefore that .NET applications can use. However note that not all .NET languages will support all the types in the CTS. The CTS is a superset of the CLS.

2.3 What is the CLS?

CLS = Common Language Specification. This is a subset of the CTS which all .NET languages are expected to support. The idea is that any program which uses CLS-compliant types can interoperate with any .NET program written in any language.

In theory this allows very tight interop between different .NET languages - for example allowing a C# class to inherit from a VB class.

2.4 What is IL?

IL = Intermediate Language. Also known as MSIL (Microsoft Intermediate Language) or CIL (Common Intermediate Language). All .NET source code (of any language) is compiled to IL. The IL is then converted to machine code at the point where the software is installed, or at run-time by a Just-In-Time (JIT) compiler.

2.5 What is C#?

C# is a new language designed by Microsoft to work with the .NET framework. In their "Introduction to C#" whitepaper, Microsoft describe C# as follows:

"C# is a simple, modern, object oriented, and type-safe programming language derived from C and C++. C# (pronounced "C sharp") is firmly planted in the C and C++ family tree of languages, and will immediately be familiar to C and C++ programmers. C# aims to combine the high productivity of Visual Basic and the raw power of C++."

Substitute 'Java' for 'C#' in the quote above, and you'll see that the statement still works pretty well :-).

If you are a C++ programmer, you might like to check out my [C# FAQ](#).

2.6 What does 'managed' mean in the .NET context?

The term 'managed' is the cause of much confusion. It is used in various places within .NET, meaning slightly different things.

Managed code: The .NET framework provides several core run-time services to the programs that run within it - for example exception handling and security. For these services to work, the code must provide a minimum level of information to the runtime. Such code is called *managed code*. All C# and Visual Basic.NET code is managed by default. VS7 C++ code is *not* managed

by default, but the compiler can produce managed code by specifying a command-line switch (/com+).

Managed *data*: This is data that is allocated and de-allocated by the .NET runtime's garbage collector. C# and VB.NET data is always managed. VS7 C++ data is unmanaged by default, even when using the /com+ switch, but it can be marked as managed using the `__gc` keyword.

Managed *classes*: This is usually referred to in the context of Managed Extensions (ME) for C++. When using ME C++, a class can be marked with the `__gc` keyword. As the name suggests, this means that the memory for instances of the class is managed by the garbage collector, but it also means more than that. The class becomes a fully paid-up member of the .NET community with the benefits and restrictions that brings. An example of a benefit is proper interop with classes written in other languages - for example, a managed C++ class can inherit from a VB class. An example of a restriction is that a managed class can only inherit from one base class.

2.7 What is reflection?

All .NET compilers produce metadata about the types defined in the modules they produce. This metadata is packaged along with the module (modules in turn are packaged together in assemblies), and can be accessed by a mechanism called **reflection**. The System.Reflection namespace contains classes that can be used to interrogate the types for a module/assembly.

Using reflection to access .NET metadata is very similar to using ITypeLib/ITypeInfo to access type library data in COM, and it is used for similar purposes - e.g. determining data type sizes for marshaling data across context/process/machine boundaries.

Reflection can also be used to dynamically invoke methods (see System.Type.InvokeMember), or even create types dynamically at run-time (see System.Reflection.Emit.TypeBuilder).

3. Assemblies

3.1 What is an assembly?

An assembly is sometimes described as a logical .EXE or .DLL, and can be an *application* (with a main entry point) or a *library*. An assembly consists of one or more files (dlls, exes, html files etc), and represents a group of resources, type definitions, and implementations of those types. An assembly may also contain references to other assemblies. These resources, types and references are described in a block of data called a *manifest*. The manifest is part of the assembly, thus making the assembly self-describing.

An important aspect of assemblies is that they are part of the identity of a type. The identity of a type is the assembly that houses it combined with the

type name. This means, for example, that if assembly A exports a type called T, and assembly B exports a type called T, the .NET runtime sees these as two completely different types. Furthermore, don't get confused between assemblies and namespaces - namespaces are merely a hierarchical way of organising type names. To the runtime, type names are type names, regardless of whether namespaces are used to organise the names. It's the assembly plus the typename (regardless of whether the type name belongs to a namespace) that uniquely identifies a type to the runtime.

Assemblies are also important in .NET with respect to security - many of the security restrictions are enforced at the assembly boundary.

Finally, assemblies are the unit of versioning in .NET - more on this below.

3.2 How can I produce an assembly?

The simplest way to produce an assembly is directly from a .NET compiler. For example, the following C# program:

```
public class CTest
{
    public CTest()
    {
        System.Console.WriteLine( "Hello from CTest" );
    }
}
```

can be compiled into a library assembly (dll) like this:

```
csc /t:library ctest.cs
```

You can then view the contents of the assembly by running the "IL Disassembler" tool that comes with the .NET SDK.

Alternatively you can compile your source into **modules**, and then combine the modules into an assembly using the assembly linker (al.exe). For the C# compiler, the /target:module switch is used to generate a module instead of an assembly.

3.3 What is the difference between a private assembly and a shared assembly?

- **Location and visibility:** A private assembly is normally used by a single application, and is stored in the application's directory, or a sub-directory beneath. A shared assembly is normally stored in the global assembly cache, which is a repository of assemblies maintained by the .NET runtime. Shared assemblies are usually libraries of code which many applications will find useful, e.g. the .NET framework classes.
- **Versioning:** The runtime enforces versioning constraints only on shared assemblies, not on private assemblies.

3.4 How do assemblies find each other?

By searching directory paths. There are several factors which can affect the path (such as the AppDomain host, and application configuration files), but for private assemblies the search path is normally the application's directory and its sub-directories. For shared assemblies, the search path is normally same as the private assembly path plus the shared assembly cache.

3.5 How does assembly versioning work?

Each assembly has a version number called the *compatibility* version. Also each reference to an assembly (from another assembly) includes both the name and version of the referenced assembly.

The version number has four numeric parts (e.g. 5.5.2.33). Assemblies with either of the first two parts different are normally viewed as incompatible. If the first two parts are the same, but the third is different, the assemblies are deemed as 'maybe compatible'. If only the fourth part is different, the assemblies are deemed compatible. However, this is just the default guideline - it is the *version policy* that decides to what extent these rules are enforced. The version policy can be specified via the application configuration file.

Remember: versioning is only applied to shared assemblies, not private assemblies.

4. Application Domains

4.1 What is an Application Domain?

An AppDomain can be thought of as a lightweight process. Multiple AppDomains can exist inside a Win32 process. The primary purpose of the AppDomain is to isolate an application from other applications.

Win32 processes provide isolation by having distinct memory address spaces. This is effective, but it is expensive and doesn't scale well. The .NET runtime enforces AppDomain isolation by keeping control over the use of memory - all memory in the AppDomain is managed by the .NET runtime, so the runtime can ensure that AppDomains do not access each other's memory.

4.2 How does an AppDomain get created?

AppDomains are usually created by *hosts*. Examples of hosts are the Windows Shell, ASP.NET and IE. When you run a .NET application from the command-line, the host is the Shell. The Shell creates a new AppDomain for every application.

AppDomains can also be explicitly created by .NET applications. Here is a C# sample which creates an AppDomain, creates an instance of an object inside

it, and then executes one of the object's methods. Note that you must name the executable 'appdomaintest.exe' for this code to work as-is.

```
using System;
using System.Runtime.Remoting;

public class CAppDomainInfo : MarshalByRefObject
{
    public string GetAppDomainInfo()
    {
        return "AppDomain = " + AppDomain.CurrentDomain.FriendlyName;
    }
}

public class App
{
    public static int Main()
    {
        AppDomain ad = AppDomain.CreateDomain( "Andy's new domain", null, null
);
        ObjectHandle oh = ad.CreateInstance( "appdomaintest",
"CAppDomainInfo" );
        CAppDomainInfo adInfo = (CAppDomainInfo)(oh.Unwrap());
        string info = adInfo.GetAppDomainInfo();

        Console.WriteLine( "AppDomain info: " + info );
        return 0;
    }
}
```

4.3 Can I write my own .NET host?

Yes. For an example of how to do this, take a look at the source for the dm.net moniker developed by Jason Whittington and Don Box (<http://staff.develop.com/jasonw/clr/readme.htm>). There is also a code sample in the .NET SDK called CorHost.

5. Garbage Collection

5.1 What is garbage collection?

Garbage collection is a system whereby a run-time component takes responsibility for managing the lifetime of objects and the heap memory that they occupy. This concept is not new to .NET - Java and many other languages/runtimes have used garbage collection for some time.

5.2 Is it true that objects don't always get destroyed immediately when the last reference goes away?

Yes. The garbage collector offers no guarantees about the time when an object will be destroyed and its memory reclaimed.

There is an interesting thread in the archives, started by Chris Sells, about the implications of non-deterministic destruction of objects in C#:

<http://discuss.develop.com/archives/wa.exe?A2=ind0007&L=DOTNET&P=R24819>

In October 2000, Microsoft's Brian Harry posted a lengthy analysis of the problem:

<http://discuss.develop.com/archives/wa.exe?A2=ind0010A&L=DOTNET&P=R28572>

Chris Sells' response to Brian's posting is here:

<http://discuss.develop.com/archives/wa.exe?A2=ind0010C&L=DOTNET&P=R983>

5.3 Why doesn't the .NET runtime offer deterministic destruction?

Because of the garbage collection algorithm. The .NET garbage collector works by periodically running through a list of all the objects that are currently being referenced by an application. All the objects that it *doesn't* find during this search are ready to be destroyed and the memory reclaimed. The implication of this algorithm is that the runtime doesn't get notified immediately when the final reference on an object goes away - it only finds out during the next sweep of the heap.

Futhermore, this type of algorithm works best by performing the garbage collection sweep as rarely as possible. Normally heap exhaustion is the trigger for a collection sweep.

5.4 Is the lack of deterministic destruction in .NET a problem?

It's certainly an issue that affects component design. If you have objects that maintain expensive or scarce resources (e.g. database locks), you need to provide some way for the client to tell the object to release the resource when it is done. Microsoft recommend that you provide a method called `Dispose()` for this purpose. However, this causes problems for distributed objects - in a distributed system who calls the `Dispose()` method? Some form of reference-counting or ownership-management mechanism is needed to handle distributed objects - unfortunately the runtime offers no help with this.

5.5 Does non-deterministic destruction affect the usage of COM objects from managed code?

Yes. When using a COM object from managed code, you are effectively relying on the garbage collector to call the final release on your object. If your COM object holds onto an expensive resource which is only cleaned-up after the final release, you may need to provide a new interface on your object which supports an explicit `Dispose()` method.

5.6 I've heard that Finalize methods should be avoided. Should I implement Finalize on my class?

An object with a Finalize method is more work for the garbage collector than an object without one. Also there are no guarantees about the order in which objects are Finalized, so there are issues surrounding access to other objects from the Finalize method. Finally, there is no guarantee that a Finalize method will get called on an object, so it should never be relied upon to do clean-up of an object's resources.

Microsoft recommend the following pattern:

```
public class CTest : IDisposable
{
    public void Dispose()
    {
        ... // Cleanup activities
        GC.SuppressFinalize(this);
    }

    ~CTest()// C# syntax hiding the Finalize() method
    {
        Dispose();
    }
}
```

In the normal case the client calls Dispose(), the object's resources are freed, and the garbage collector is relieved of its Finalizing duties by the call to SuppressFinalize(). In the worst case, i.e. the client forgets to call Dispose(), there is a reasonable chance that the object's resources will eventually get freed by the garbage collector calling Finalize(). Given the limitations of the garbage collection algorithm this seems like a pretty reasonable approach.

5.7 Do I have any control over the garbage collection algorithm?

A little. For example, the System.GC class exposes a Collect method - this forces the garbage collector to collect all unreferenced objects immediately.

5.8 How can I find out what the garbage collector is doing?

Lots of interesting statistics are exported from the .NET runtime via the '.NET CLR xxx' performance counters. Use Performance Monitor to view them.

6. Serialization

6.1 What is serialization?

Serialization is the process of converting an object into a stream of bytes. Deserialization is the opposite process of creating an object from a stream of

bytes. Serialization/Deserialization is mostly used to transport objects (e.g. during remoting), or to persist objects (e.g. to a file or database).

6.2 Does the .NET Framework have in-built support for serialization?

There are two separate mechanisms provided by the .NET class library - XmlSerializer and SoapFormatter/BinaryFormatter. Microsoft uses XmlSerializer for Web Services, and uses SoapFormatter/BinaryFormatter for remoting. Both are available for use in your own code.

6.3 I want to serialize instances of my class. Should I use XmlSerializer, SoapFormatter or BinaryFormatter?

It depends. XmlSerializer has severe limitations such as the requirement that the target class has a parameterless constructor, and only public read/write properties and fields can be serialized. However, on the plus side, XmlSerializer has good support for customising the XML document that is produced or consumed. XmlSerializer's features mean that it is most suitable for cross-platform work, or for constructing objects from existing XML documents.

SoapFormatter and BinaryFormatter have fewer limitations than XmlSerializer. They can serialize private fields, for example. However they both require that the target class be marked with the [Serializable] attribute, so like XmlSerializer the class needs to be written with serialization in mind. Also there are some quirks to watch out for - for example on deserialization the constructor of the new object is not invoked.

The choice between SoapFormatter and BinaryFormatter depends on the application. BinaryFormatter makes sense where both serialization and deserialization will be performed on the .NET platform and where performance is important. SoapFormatter generally makes more sense in all other cases, for ease of debugging if nothing else.

6.4 Can I customise the serialization process?

Yes. XmlSerializer supports a range of attributes that can be used to configure serialization for a particular class. For example, a field or property can be marked with the [XmlIgnore] attribute to exclude it from serialization. Another example is the [XmlElement] attribute, which can be used to specify the XML element name to be used for a particular property or field.

Serialization via SoapFormatter/BinaryFormatter can also be controlled to some extent by attributes. For example, the [NonSerialized] attribute is the equivalent of XmlSerializer's [XmlIgnore] attribute. Ultimate control of the serialization process can be achieved by implementing the the ISerializable interface on the class whose instances are to be serialized.

6.5 Why is XmlSerializer so slow?

There is a once-per-process-per-type overhead with XmlSerializer. So the first time you serialize or deserialize an object of a given type in an application, there is a significant delay. This normally doesn't matter, but it may mean, for example, that XmlSerializer is a poor choice for loading configuration settings during startup of a GUI application.

6.6 Why do I get errors when I try to serialize a Hashtable?

XmlSerializer will refuse to serialize instances of any class that implements IDictionary, e.g. Hashtable. SoapFormatter and BinaryFormatter do not have this restriction.

6.7 XmlSerializer is throwing a generic "There was an error reflecting MyClass" error. How do I find out what the problem is?

Look at the InnerException property of the exception that is thrown to get a more specific error message.

7. Attributes

7.1 What are attributes?

There are at least two types of .NET attribute. The first type I will refer to as a metadata attribute - it allows some data to be attached to a class or method. This data becomes part of the metadata for the class, and (like other class metadata) can be accessed via reflection. An example of a metadata attribute is [serializable], which can be attached to a class and means that instances of the class can be serialized.

```
[serializable] public class CTest {}
```

The other type of attribute is a context attribute. Context attributes use a similar syntax to metadata attributes but they are fundamentally different. Context attributes provide an interception mechanism whereby instance activation and method calls can be pre- and/or post-processed. If you've come across Keith Brown's universal delegator you'll be familiar with this idea.

7.2 Can I create my own metadata attributes?

Yes. Simply derive a class from System.Attribute and mark it with the AttributeUsage attribute. For example:

```
[AttributeUsage(AttributeTargets.Class)]
public class InspiredByAttribute : System.Attribute
{
    public string InspiredBy;

    public InspiredByAttribute( string inspiredBy )
    {
```

```

        InspiredBy = inspiredBy;
    }
}

[InspiredBy("Andy Mc's brilliant .NET FAQ")]
class CTest
{
}

class CApp
{
    public static void Main()
    {
        object[] atts = typeof(CTest).GetCustomAttributes(true);

        foreach( object att in atts )
            if( att is InspiredByAttribute )
                Console.WriteLine( "Class CTest was inspired by {0}",
((InspiredByAttribute)att).InspiredBy );
    }
}

```

7.3 Can I create my own context attributes?

Yes. Take a look at Don Box's sample (called CallThreshold) at <http://www.develop.com/dbox/dotnet/threshold/>, and also Peter Drayton's Tracehook.NET at <http://www.razorsoft.net/>

8. Code Access Security

8.1 What is Code Access Security (CAS)?

CAS is the part of the .NET security model that determines whether or not a piece of code is allowed to run, and what resources it can use when it is running. For example, it is CAS that will prevent a .NET web applet from formatting your hard disk.

8.2 How does CAS work?

The CAS security policy revolves around two key concepts - code groups and permissions. Each .NET assembly is a member of a particular **code group**, and each code group is granted the permissions specified in a **named permission set**.

For example, using the default security policy, a control downloaded from a web site belongs to the 'Zone - Internet' code group, which adheres to the permissions defined by the 'Internet' named permission set. (Naturally the 'Internet' named permission set represents a very restrictive range of permissions.)

8.3 Who defines the CAS code groups?

Microsoft defines some default ones, but you can modify these and even create your own. To see the code groups defined on your system, run 'caspol -lg' from the command-line. On my system it looks like this:

Level = Machine

Code Groups:

```
1. All code: Nothing
  1.1. Zone - MyComputer: FullTrust
    1.1.1. Honor SkipVerification requests: SkipVerification
  1.2. Zone - Intranet: LocalIntranet
  1.3. Zone - Internet: Internet
  1.4. Zone - Untrusted: Nothing
  1.5. Zone - Trusted: Internet
  1.6. StrongName -
0024000004800000940000000602000000240000525341310004000003
000000CFCB3291AA715FE99D40D49040336F9056D7886FED46775BC7BB5430BA4444FEF834
8EBD06
F962F39776AE4DC3B7B04A7FE6F49F25F740423EBF2C0B89698D8D08AC48D69CED0FC8F83B
465E08
07AC11EC1DCC7D054E807A43336DDE408A5393A48556123272CEEEE72F1660B71927D3856
1AABF5C
AC1DF1734633C602F8F2D5: Everything
```

Note the hierarchy of code groups - the top of the hierarchy is the most general ('All code'), which is then sub-divided into several groups, each of which in turn can be sub-divided. Also note that (somewhat counter-intuitively) a sub-group can be associated with a more permissive permission set than its parent.

8.4 How do I define my own code group?

Use caspol. For example, suppose you trust code from www.mydomain.com and you want it have full access to your system, but you want to keep the default restrictions for all other internet sites. To achieve this, you would add a new code group as a sub-group of the 'Zone - Internet' group, like this:

```
caspol -ag 1.3 -site www.mydomain.com FullTrust
```

Now if you run caspol -lg you will see that the new group has been added as group 1.3.1:

```
...
  1.3. Zone - Internet: Internet
    1.3.1. Site - www.mydomain.com: FullTrust
  ...
```

Note that the numeric label (1.3.1) is just a caspol invention to make the code groups easy to manipulate from the command-line. The underlying runtime never sees it.

8.5 How do I change the permission set for a code group?

Use caspol. If you are the machine administrator, you can operate at the 'machine' level - which means not only that the changes you make become the default for the machine, but also that users cannot change the permissions to be more permissive. If you are a normal (non-admin) user you can still modify the permissions, but only to make them more restrictive. For example, to allow intranet code to do what it likes you might do this:

```
caspol -cg 1.2 FullTrust
```

Note that because this is more permissive than the default policy (on a standard system), you should only do this at the machine level - doing it at the user level will have no effect.

8.6 Can I create my own permission set?

Yes. Use caspol -ap, specifying an XML file containing the permissions in the permission set. To save you some time, [here](#) is a sample file corresponding to the 'Everything' permission set - just edit to suit your needs. When you have edited the sample, add it to the range of available permission sets like this:

```
caspol -ap samplepermset.xml
```

Then, to apply the permission set to a code group, do something like this:

```
caspol -cg 1.3 SamplePermSet
```

(By default, 1.3 is the 'Internet' code group)

8.7 I'm having some trouble with CAS. How can I diagnose my problem?

Caspol has a couple of options that might help. First, you can ask caspol to tell you what code group an assembly belongs to, using caspol -rsg. Similarly, you can ask what permissions are being applied to a particular assembly using caspol -rsp.

8.8 I can't be bothered with all this CAS stuff. Can I turn it off?

Yes, as long as you are an administrator. Just run:

```
caspol -s off
```

9. Intermediate Language (IL)

9.1 Can I look at the IL for an assembly?

Yes. MS supply a tool called Ildasm which can be used to view the metadata and IL for an assembly.

9.2 Can source code be reverse-engineered from IL?

Yes, it is often relatively straightforward to regenerate high-level source (e.g. C#) from IL.

9.3 How can I stop my code being reverse-engineered from IL?

There is currently no simple way to stop code being reverse-engineered from IL. In future it is likely that IL obfuscation tools will become available, either from MS or from third parties. These tools work by 'optimising' the IL in such a way that reverse-engineering becomes much more difficult.

Of course if you are writing web services then reverse-engineering is not a problem as clients do not have access to your IL.

9.4 Can I write IL programs directly?

Yes. [Peter Drayton](#) posted this simple example to the DOTNET mailing list:

```
.assembly MyAssembly {}
.class MyApp {
  .method static void Main() {
    .entrypoint
    ldstr    "Hello, IL!"
    call    void System.Console::WriteLine(class System.Object)
    ret
  }
}
```

Just put this into a file called hello.il, and then run `ilasm hello.il`. An exe assembly will be generated.

9.5 Can I do things in IL that I can't do in C#?

Yes. A couple of simple examples are that you can throw exceptions that are not derived from `System.Exception`, and you can have non-zero-based arrays.

10. Implications for COM

10.1 Is COM dead?

This subject causes a lot of controversy, as you'll see if you read the mailing list archives. Take a look at the following two threads:

<http://discuss.develop.com/archives/wa.exe?A2=ind0007&L=DOTNET&D=0&P=68241>

<http://discuss.develop.com/archives/wa.exe?A2=ind0007&L=DOTNET&P=R60761>

FWIW my view is as follows: COM is many things, and it's different things to different people. But to me, COM is fundamentally about how little blobs of code find other little blobs of code, and how they communicate with each other when they find each other. COM specifies precisely how this location and communication takes place. In a 'pure' .NET world, consisting entirely of .NET objects, little blobs of code still find each other and talk to each other, but they don't use COM to do so. They use a model which is similar to COM in some ways - for example, type information is stored in a tabular form packaged with the component, which is quite similar to packaging a type library with a COM component. But it's not COM.

So, does this matter? Well, I don't really care about most of the COM stuff going away - I don't care that finding components doesn't involve a trip to the registry, or that I don't use IDL to define my interfaces. But there is one thing that I wouldn't like to go away - I wouldn't like to lose the idea of interface-based development. COM's greatest strength, in my opinion, is its insistence on a cast-iron separation between interface and implementation. Unfortunately, the .NET framework seems to make no such insistence - it lets you do interface-based development, but it doesn't insist. Some people would argue that having a choice can never be a bad thing, and maybe they're right, but I can't help feeling that maybe it's a backward step.

10.2 Is DCOM dead?

Pretty much, for .NET developers. The .NET Framework has a new remoting model which is not based on DCOM. Of course DCOM will still be used in interop scenarios.

10.3 Is MTS/COM+ dead?

No. The approach for the first .NET release is to provide access to the existing COM+ services (through an interop layer) rather than replace the services with native .NET ones. Various tools and attributes are provided to try to make this as painless as possible. The PDC release of the .NET SDK includes interop support for core services (JIT activation, transactions) but not some of the higher level services (e.g. COM+ Events, Queued components).

Over time it is expected that interop will become more seamless - this may mean that some services become a core part of the CLR, and/or it may mean that some services will be rewritten as managed code which runs on top of the CLR.

For more on this topic, search for postings by Joe Long in the archives - Joe is the MS group manager for COM+. Start with this message:

<http://discuss.develop.com/archives/wa.exe?A2=ind0007&L=DOTNET&P=R68370>

10.4 Can I use COM components from .NET programs?

Yes. COM components are accessed from the .NET runtime via a Runtime Callable Wrapper (RCW). This wrapper turns the COM interfaces exposed by the COM component into .NET-compatible interfaces. For oleautomation interfaces, the RCW can be generated automatically from a type library. For non-oleautomation interfaces, it may be necessary to develop a custom RCW which manually maps the types exposed by the COM interface to .NET-compatible types.

Here's a simple example for those familiar with ATL. First, create an ATL component which implements the following IDL:

```
import "oaidl.idl";
import "ocidl.idl";

[
    object,
    uuid(EA013F93-487A-4403-86EC-FD9FEE5E6206),
    helpstring("ICppName Interface"),
    pointer_default(unique),
    oleautomation
]

interface ICppName : IUnknown
{
    [helpstring("method SetName")] HRESULT SetName([in] BSTR name);
    [helpstring("method GetName")] HRESULT GetName([out,retval] BSTR *pName );
};

[
    uuid(F5E4C61D-D93A-4295-A4B4-2453D4A4484D),
    version(1.0),
    helpstring("cppcomserver 1.0 Type Library")
]
library CPPCOMSERVERLib
{
    importlib("stdole32.tlb");
    importlib("stdole2.tlb");
    [
        uuid(600CE6D9-5ED7-4B4D-BB49-E8D5D5096F70),
        helpstring("CppName Class")
    ]
    coclass CppName
    {
        [default] interface ICppName;
    };
};
```

When you've built the component, you should get a typelibrary. Run the TLBIMP utility on the typelibrary, like this:

```
tlbimp cppcomserver.tlb
```

If successful, you will get a message like this:

```
Typelib imported successfully to CPPCOMSERVERLib.dll
```

You now need a .NET client - let's use C#. Create a .cs file containing the following code:

```
using System;
using CPPCOMSERVERLib;

public class MainApp
{
    static public void Main()
    {
        CppName cppname = new CppName();
        cppname.SetName( "bob" );
        Console.WriteLine( "Name is " + cppname.GetName() );
    }
}
```

Note that we are using the type library name as a namespace, and the COM class name as the class. Alternatively we could have used CPPCOMSERVERLib.CppName for the class name and gone without the using CPPCOMSERVERLib statement.

Compile the C# code like this:

```
csc /r:cppcomserverlib.dll csharpcomclient.cs
```

Note that the compiler is being told to reference the DLL we previously generated from the typelibrary using TLBIMP.

You should now be able to run csharpcomclient.exe, and get the following output on the console:

```
Name is bob
```

10.5 Can I use .NET components from COM programs?

Yes. .NET components are accessed from COM via a COM Callable Wrapper (CCW). This is similar to a RCW (see previous question), but works in the opposite direction. Again, if the wrapper cannot be automatically generated by the .NET development tools, or if the automatic behaviour is not desirable, a custom CCW can be developed. Also, for COM to 'see' the .NET component, the .NET component must be registered in the registry.

Here's a simple example. Create a C# file called testcomserver.cs and put the following in it:

```
using System;
using System.Runtime.InteropServices;

namespace AndyMc
{
    [ClassInterface(ClassInterfaceType.AutoDual)]
    public class CSharpCOMServer
```

```

    {
        public CSharpCOMServer() {}
        public void SetName( string name ) { m_name = name; }
        public string GetName() { return m_name; }
        private string m_name;
    }
}

```

Then compile the .cs file as follows:

```
csc /target:library testcomserver.cs
```

You should get a dll, which you register like this:

```
regasm testcomserver.dll /tlb:testcomserver.tlb /codebase
```

Now you need to create a client to test your .NET COM component. VBScript will do - put the following in a file called comclient.vbs:

```

Dim dotNetObj
Set dotNetObj = CreateObject("AndyMc.CSharpCOMServer")
dotNetObj.SetName ("bob")
MsgBox "Name is " & dotNetObj.GetName()

```

and run the script like this:

```
wscript comclient.vbs
```

And hey presto you should get a message box displayed with the text "Name is bob".

An alternative to the approach above it to use the dm.net moniker developed by Jason Whittington and Don Box. Go to <http://staff.develop.com/jasonw/clr/readme.htm> to check it out.

10.6 Is ATL redundant in the .NET world?

Yes, if you are writing applications that live inside the .NET framework. Of course many developers may wish to continue using ATL to write C++ COM components that live outside the framework, but if you are inside you will almost certainly want to use C#. Raw C++ (and therefore ATL which is based on it) doesn't have much of a place in the .NET world - it's just too near the metal and provides too much flexibility for the runtime to be able to manage it.

11. Miscellaneous

11.1 How does .NET remoting work?

.NET remoting involves sending messages along channels. Two of the standard channels are HTTP and TCP. TCP is intended for LANs only - HTTP can be used for LANs or WANs (internet).

Support is provided for multiple message serialization formats. Examples are SOAP (XML-based) and binary. By default, the HTTP channel uses SOAP (via the .NET runtime Serialization SOAP Formatter), and the TCP channel uses binary (via the .NET runtime Serialization Binary Formatter). But either channel can use either serialization format.

There are a number of styles of remote access:

- *SingleCall*. Each incoming request from a client is serviced by a new object. The object is thrown away when the request has finished.
- *Singleton*. All incoming requests from clients are processed by a single server object.
- *Client-activated object*. This is the old stateful (D)COM model whereby the client receives a reference to the remote object and holds that reference (thus keeping the remote object alive) until it is finished with it.

Distributed garbage collection of objects is managed by a system called 'leased based lifetime'. Each object has a lease time, and when that time expires the object is disconnected from the .NET runtime remoting infrastructure. Objects have a default renew time - the lease is renewed when a successful call is made from the client to the object. The client can also explicitly renew the lease.

If you're interested in using XML-RPC as an alternative to SOAP, take a look at Charles Cook's XML-RPC.Net site at <http://www.cookcomputing.com/xmlrpc/xmlrpc.shtml>.

11.2 How can I get at the Win32 API from a .NET program?

Use P/Invoke. This uses similar technology to COM Interop, but is used to access static DLL entry points instead of COM objects. Here is an example of C# calling the Win32 MessageBox function:

```
using System;
using System.Runtime.InteropServices;

class MainApp
{
    [DllImport("user32.dll", EntryPoint="MessageBox", SetLastError=true,
    CharSet=CharSet.Auto)]
    public static extern int MessageBox(int hWnd, String strMessage, String strCaption,
    uint uiType);

    public static void Main()
```

```
    {  
        MessageBox( 0, "Hello, this is PInvoke in operation!", ".NET", 0 );  
    }  
}
```

12. Class Library

12.1 File I/O

12.1.1 How do I read from a text file?

First, use a `System.IO.FileStream` object to open the file:

```
FileStream fs = new FileStream( @"c:\test.txt", FileMode.Open, FileAccess.Read );
```

`FileStream` inherits from `Stream`, so you can wrap the `FileStream` object with a `StreamReader` object. This provides a nice interface for processing the stream line by line:

```
StreamReader sr = new StreamReader( fs );  
string curLine;  
while( (curLine = sr.ReadLine()) != null )  
    Console.WriteLine( curLine );
```

Finally close the `StreamReader` object:

```
sr.Close();
```

Note that this will automatically call `Close()` on the underlying `Stream` object, so an explicit `fs.Close()` is not required.

12.1.2 How do I write to a text file?

Similar to the read example, except use `StreamWriter` instead of `StreamReader`.

12.1.3 How do I read/write binary files?

Similar to text files, except wrap the `FileStream` object with a `BinaryReader/Writer` object instead of a `StreamReader/Writer` object.

12.2 Text Processing

12.2.1 Are regular expressions supported?

Yes. Use the `System.Text.RegularExpressions.Regex` class. For example, the following code updates the title in an HTML file:

```
FileStream fs = new FileStream( "test.htm", FileMode.Open, FileAccess.Read );  
StreamReader sr = new StreamReader( fs );  
  
Regex r = new Regex( "<TITLE>(.*?)</TITLE>" );
```

```

string s;
while( (s = sr.ReadLine()) != null )
{
    if( r.IsMatch( s ) )
        s = r.Replace( s, "<TITLE>New and improved ${1}</TITLE>" );
    Console.WriteLine( s );
}

```

12.3 Internet

12.3.1 How do I download a web page?

First use the `System.Net.WebRequestFactory` class to acquire a `WebRequest` object:

```
WebRequest request = WebRequest.Create( "http://localhost" );
```

Then ask for the response from the request:

```
WebResponse response = request.GetResponse();
```

The `GetResponse` method blocks until the download is complete. Then you can access the response stream like this:

```
Stream s = response.GetResponseStream();

// Output the downloaded stream to the console
StreamReader sr = new StreamReader( s );
string line;
while( (line = sr.ReadLine()) != null )
    Console.WriteLine( line );
```

Note that `WebRequest` and `WebResponse` objects can be downcast to `HttpWebRequest` and `HttpWebResponse` objects respectively, to access http-specific functionality.

12.3.2 How do I use a proxy?

Two approaches - to affect all web requests do this:

```
System.Net.GlobalProxySelection.Select = new WebProxy( "proxyname", 80 );
```

Alternatively, to set the proxy for a specific web request, do this:

```
HttpWebRequest request = (HttpWebRequest)WebRequest.Create( "http://localhost" );
request.Proxy = new WebProxy( "proxyname", 80 );
```

12.4 XML

12.4.1 Is DOM supported?

Yes. Take this example XML document:

```
<PEOPLE>
  <PERSON>Fred</PERSON>
  <PERSON>Bill</PERSON>
</PEOPLE>
```

This document can be parsed as follows:

```
XmlDocument doc = new XmlDocument();
doc.Load( "test.xml" );

XmlNode root = doc.DocumentElement;

foreach( XmlNode personElement in root.ChildNodes )
    Console.WriteLine( personElement.FirstChild.Value.ToString() );
```

The output is:

```
Fred
Bill
```

12.4.2 Is SAX supported?

No. Instead, a new XmlReader/XmlWriter API is offered. Like SAX it is stream-based but it uses a 'pull' model rather than SAX's 'push' model. Here's an example:

```
XmlTextReader reader = new XmlTextReader( "test.xml" );

while( reader.Read() )
{
    if( reader.NodeType == XmlNodeType.Element && reader.Name == "PERSON" )
    {
        reader.Read(); // Skip to the child text
        Console.WriteLine( reader.Value );
    }
}
```

12.4.3 Is XPath supported?

Yes, via the XPathXXX classes:

```
XPathDocument xpdoc = new XPathDocument("test.xml");
XPathNavigator nav = xpdoc.CreateNavigator();
XPathExpression expr = nav.Compile("descendant::PEOPLE/PERSON");

XPathNodeIterator iterator = nav.Select(expr);
while (iterator.MoveNext())
    Console.WriteLine(iterator.Current);
```

12.5 Threading

12.5.1 Is multi-threading supported?

Yes, there is extensive support for multi-threading. New threads can be spawned, and there is a system-provided threadpool which applications can use.

12.5.2 How do I spawn a thread?

Create an instance of a `System.Threading.Thread` object, passing it an instance of a `ThreadStart` delegate that will be executed on the new thread. For example:

```
class MyThread
{
    public MyThread( string initData )
    {
        m_data = initData;
        m_thread = new Thread( new ThreadStart(ThreadMain) );
        m_thread.Start();
    }

    // ThreadMain() is executed on the new thread.
    private void ThreadMain()
    {
        Console.WriteLine( m_data );
    }

    public void WaitUntilFinished()
    {
        m_thread.Join();
    }

    private Thread m_thread;
    private string m_data;
}
```

In this case creating an instance of the `MyThread` class is sufficient to spawn the thread and execute the `MyThread.ThreadMain()` method:

```
MyThread t = new MyThread( "Hello, world." );
t.WaitUntilFinished();
```

12.5.3 How do I stop a thread?

There are several options. First, you can use your own communication mechanism to tell the `ThreadStart` method to finish. Alternatively the `Thread` class has in-built support for instructing the thread to stop. The two principle methods are `Thread.Interrupt()` and `Thread.Abort()`. The former will cause a `ThreadInterruptedException` to be thrown on the thread when it next goes into a `WaitJoinSleep` state. In other words, `Thread.Interrupt` is a polite way of asking the thread to stop when it is no longer doing any useful work. In contrast, `Thread.Abort()` throws a `ThreadAbortException` regardless of what the thread is doing. Furthermore, the `ThreadAbortException` cannot normally be caught (though the `ThreadStart`'s `finally` method **will** be executed). `Thread.Abort()` is a heavy-handed mechanism which should not normally be required.

12.5.4 How do I use the thread pool?

By passing an instance of a WaitCallback delegate to the ThreadPool.QueueUserWorkItem() method:

```
class CApp
{
    static void Main()
    {
        string s = "Hello, World";
        ThreadPool.QueueUserWorkItem( new WaitCallback( DoWork ), s );

        Thread.Sleep( 1000 );    // Give time for work item to be executed
    }

    // DoWork is executed on a thread from the thread pool.
    static void DoWork( object state )
    {
        Console.WriteLine( state );
    }
}
```

12.5.5 How do I know when my thread pool work item has completed?

There is no way to query the thread pool for this information. You must put code into the WaitCallback method to signal that it has completed. Events are useful for this.

12.5.6 How do I prevent concurrent access to my data?

Each object has a concurrency lock (critical section) associated with it. The System.Threading.Monitor.Enter/Exit methods are used to acquire and release this lock. For example, instances of the following class only allow one thread at a time to enter method f():

```
class C
{
    public void f()
    {
        try
        {
            Monitor.Enter(this);
            ...
        }
        finally
        {
            Monitor.Exit(this);
        }
    }
}
```

C# has a 'lock' keyword which provides a convenient shorthand for the code above:

```
class C
```

```

{
    public void f()
    {
        lock(this)
        {
            ...
        }
    }
}

```

Note that calling `Monitor.Enter(myObject)` does NOT mean that all access to `myObject` is serialized. It means that the synchronisation lock associated with `myObject` has been acquired, and no other thread can acquire that lock until `Monitor.Exit(o)` is called. In other words, this class is functionally equivalent to the classes above:

```

class C
{
    public void f()
    {
        lock( m_object )
        {
            ...
        }
    }

    private m_object = new object();
}

```

12.6 Tracing

12.6.1 Is there built-in support for tracing/logging?

Yes, in the `System.Diagnostics` namespace. There are two main classes that deal with tracing - `Debug` and `Trace`. They both work in a similar way - the difference is that tracing from the `Debug` class only works in builds that have the `DEBUG` symbol defined, whereas tracing from the `Trace` class only works in builds that have the `TRACE` symbol defined. Typically this means that you should use `System.Diagnostics.Trace.WriteLine` for tracing that you want to work in debug and release builds, and `System.Diagnostics.Debug.WriteLine` for tracing that you want to work only in debug builds.

12.6.2 Can I redirect tracing to a file?

Yes. The `Debug` and `Trace` classes both have a `Listeners` property, which is a collection of sinks that receive the tracing that you send via `Debug.WriteLine` and `Trace.WriteLine` respectively. By default the `Listeners` collection contains a single sink, which is an instance of the `DefaultTraceListener` class. This sends output to the `Win32 OutputDebugString()` function and also the `System.Diagnostics.Debugger.Log()` method. This is useful when debugging, but if you're trying to trace a problem at a customer site, redirecting the output to a file is more appropriate. Fortunately, the `TextWriterTraceListener` class is provided for this purpose.

Here's how to use the `TextWriterTraceListener` class to redirect Trace output to a file:

```
Trace.Listeners.Clear();
FileStream fs = new FileStream( @"c:\log.txt", FileMode.Create, FileAccess.Write );
Trace.Listeners.Add( new TextWriterTraceListener( fs ) );

Trace.WriteLine( @"This will be written to c:\log.txt!" );
Trace.Flush();
```

Note the use of `Trace.Listeners.Clear()` to remove the default listener. If you don't do this, the output will go to the file *and* `OutputDebugString()`. Typically this is not what you want, because `OutputDebugString()` imposes a big performance hit.

12.6.3 Can I customise the trace output?

Yes. You can write your own `TraceListener`-derived class, and direct all output through it. Here's a simple example, which derives from `TextWriterTraceListener` (and therefore has in-built support for writing to files, as shown above) and adds timing information and the thread ID for each trace line:

```
class MyListener : TextWriterTraceListener
{
    public MyListener( Stream s ) : base(s)
    {
    }

    public override void WriteLine( string s )
    {
        Writer.WriteLine( "{0:D8} [{1:D4}] {2}",
            Environment.TickCount - m_startTickCount,
            AppDomain.GetCurrentThreadId(),
            s );
    }

    protected int m_startTickCount = Environment.TickCount;
}
```

(Note that this implementation is not complete - the `TraceListener.WriteLine` method is not overridden for example.)

The beauty of this approach is that when an instance of `MyListener` is added to the `Trace.Listeners` collection, all calls to `Trace.WriteLine()` go through `MyListener`, including calls made by referenced assemblies that know nothing about the `MyListener` class.

13. Resources

13.1 Recommended books

I recommend the following books, either because I personally like them, or because I think they are well regarded by other .NET developers. (Note that I get a commission from Amazon if you buy a book after following one of these links.)

- [Applied Microsoft .NET Framework Programming - Jeffrey Richter](#)
Much anticipated, mainly due to Richter's superb Win32 books, and most people think it delivers. The 'applied' is a little misleading - this book is mostly about how the .NET Framework works 'under the hood'. Examples are in C#, but there is also a separate [VB edition](#) of the book.
- [Essential .NET Volume 1, The Common Language Runtime - Don Box](#)
A superb book, which I recommend to anyone who already has some .NET development experience, and wants to get a deeper understanding of CLR fundamentals. It's clear that Box has deeply researched the topics and then carefully constructed a coherent story around his findings. It's rare to find such craft in a .NET text.
- [C# and the .NET Platform, 2nd Edition - Andrew Troelsen](#)
Regarded by many as the best all round C#/.NET book. Wide coverage including Windows Forms, COM interop, ADO.NET, ASP.NET etc. Troelsen also has a respected VB.NET book called [Visual Basic .NET and the .NET Platform: An Advanced Guide](#).
- [Programming Windows with C# - Charles Petzold](#)
Another slightly misleading title - this book is solely about GUI programming - Windows Forms and GDI+. Well written, with comprehensive coverage. My only (minor) criticism is that the book sticks closely to the facts, without offering a great deal in the way of 'tips and tricks' for real-world apps.
- [Windows Forms Programming in C# - Chris Sells](#)
I haven't read this myself yet, but anything Sells writes is usually worth reading.
- [Developing Applications with Visual Studio.NET - Richard Grimes](#)
Covers lots of interesting topics that other books don't, including ATL7, Managed C++, internationalization, remoting, as well as the more run-of-the-mill CLR and C# stuff. Also a lot of info on the Visual Studio IDE. This book is most suitable for reasonably experienced C++ programmers.
- [Programming Microsoft Visual Basic .NET - Francesco Balena](#)
Balena is a renowned VB-er, and the reviews of his VB.NET book are glowing.
- [.NET and COM - The Complete Interoperability Guide - Adam Nathan](#)
Don't be put off by the size - this book is very easy to digest thanks to

the superb writing style. The bible of .NET/COM interop.

- [Advanced .NET Remoting - Ingo Rammer](#)
Widely recommended.

13.2 Internet Resources

- The Microsoft .NET homepage is at <http://www.microsoft.com/net/>. Microsoft also host [GOTDOTNET](#).
- DevX host the [.NET Zone](#).
- http://www.cetus-links.org/oo_dotnet.html is a superb set of links to .NET resources.
- Chris Sells has a great set of .NET links at <http://www.sellsbrothers.com/links/#manlinks>.
- [CSharp.org](#)
- [microsoft.public.dotnet.*](#) newsgroups
- My [C# FAQ for C++ Programmers](#).

13.3 Weblogs

The following Weblogs ('blogs') have regular .NET content:

- [The .NET Guy \(Brad Wilson\)](#)
- [Charles Cook](#): Developer of XML-RPC.NET.
- [Gwyn Cole](#): Co-author of [Developing WMI solutions](#).
- [Chris Brumme](#)
- [Brad Abrams](#)
- [Don Box](#)
- [John Lam](#)
- [Peter Drayton](#): Co-author of [C# Essentials](#) and [C# in a Nutshell](#).
- [Ingo Rammer](#): Author of [Advanced .NET remoting](#).
- [Drew Marsh](#)
- [Tomas Restrepo](#)
- [Justin Rudd](#)
- [Simon Fell](#): Developer of PocketSOAP.
- [Richard Caetano](#)
- [Chris Sells](#)

13.4 Sample code & utilities

Lutz Roeder has some great utilities and libraries at <http://www.aisto.com/roeder/dotnet/>

Peter Drayton's .NET Goodies page is at <http://www.razorsoft.net/>

Don Box & Jason Whittington's dm.net COM moniker at <http://staff.develop.com/jasonw/clr/readme.htm>

Mike Woodring has some .NET samples at
<http://staff.develop.com/woodring/dotnet/>

Charles Cook's XML-RPC.Net library is available at
<http://www.cookcomputing.com/>.

Microsoft SQL Server # Interview Questions - (last updated on)

- [Transact-SQL Optimization Tips](#)
- [Index Optimization tips](#)
- [T-SQL Queries](#)
- [Data Types](#)
- [Index](#)
- [Joins](#)
- [Lock](#)
- [Stored Procedure](#)
- [Trigger](#)
- [View](#)
- [Transaction](#)
- [Other](#)
- [XML](#)
- [Tools](#)
- [Permission](#)
- [Administration](#)

Transact-SQL Optimization Tips

- **Use views and stored procedures instead of heavy-duty queries.**
This can reduce network traffic, because your client will send to server only stored procedure or view name (perhaps with some parameters) instead of large heavy-duty queries text. This can be used to facilitate permission management also, because you can restrict user access to table columns they should not see.
- **Try to use constraints instead of triggers, whenever possible.**
Constraints are much more efficient than triggers and can boost performance. So, you should use constraints instead of triggers, whenever possible.
- **Use table variables instead of temporary tables.**
Table variables require less locking and logging resources than temporary tables, so table variables should be used whenever possible. The table variables are available in SQL Server 2000 only.
- **Try to use UNION ALL statement instead of UNION, whenever possible.**
The UNION ALL statement is much faster than UNION, because UNION ALL statement does not look for duplicate rows, and UNION statement does look for duplicate rows, whether or not they exist.
- **Try to avoid using the DISTINCT clause, whenever possible.**
Because using the DISTINCT clause will result in some performance degradation, you should use this clause only when it is necessary.

- **Try to avoid using SQL Server cursors, whenever possible.**
SQL Server cursors can result in some performance degradation in comparison with select statements. Try to use correlated sub-query or derived tables, if you need to perform row-by-row operations.
- **Try to avoid the HAVING clause, whenever possible.**
The HAVING clause is used to restrict the result set returned by the GROUP BY clause. When you use GROUP BY with the HAVING clause, the GROUP BY clause divides the rows into sets of grouped rows and aggregates their values, and then the HAVING clause eliminates undesired aggregated groups. In many cases, you can write your select statement so, that it will contain only WHERE and GROUP BY clauses without HAVING clause. This can improve the performance of your query.
- **If you need to return the total table's row count, you can use alternative way instead of SELECT COUNT(*) statement.**
Because SELECT COUNT(*) statement make a full table scan to return the total table's row count, it can take very many time for the large table. There is another way to determine the total row count in a table. You can use sysindexes system table, in this case. There is ROWS column in the sysindexes table. This column contains the total row count for each table in your database. So, you can use the following select statement instead of SELECT COUNT(*): SELECT rows FROM sysindexes WHERE id = OBJECT_ID('table_name') AND indid < 2 So, you can improve the speed of such queries in several times.
- **Include SET NOCOUNT ON statement into your stored procedures to stop the message indicating the number of rows affected by a T-SQL statement.**
This can reduce network traffic, because your client will not receive the message indicating the number of rows affected by a T-SQL statement.
- **Try to restrict the queries result set by using the WHERE clause.**
This can results in good performance benefits, because SQL Server will return to client only particular rows, not all rows from the table(s). This can reduce network traffic and boost the overall performance of the query.
- **Use the select statements with TOP keyword or the SET ROWCOUNT statement, if you need to return only the first n rows.**
This can improve performance of your queries, because the smaller result set will be returned. This can also reduce the traffic between the server and the clients.
- **Try to restrict the queries result set by returning only the particular columns from the table, not all table's columns.**
This can results in good performance benefits, because SQL Server will return to client only particular columns, not all table's columns. This can reduce network traffic and boost the overall performance of the query.

- 1.Indexes
- 2.avoid more number of triggers on the table
- 3.unnecessary complicated joins
- 4.correct use of Group by clause with the select list

5.in worst cases Denormalization

Index Optimization tips

- Every index increases the time it takes to perform INSERTS, UPDATES and DELETES, so the number of indexes should not be very much. Try to use maximum 4-5 indexes on one table, not more. If you have read-only table, then the number of indexes may be increased.
- Keep your indexes as narrow as possible. This reduces the size of the index and reduces the number of reads required to read the index.
- Try to create indexes on columns that have integer values rather than character values.
- If you create a composite (multi-column) index, the order of the columns in the key are very important. Try to order the columns in the key as to enhance selectivity, with the most selective columns to the leftmost of the key.
- If you want to join several tables, try to create surrogate integer keys for this purpose and create indexes on their columns.
- Create surrogate integer primary key (identity for example) if your table will not have many insert operations.
- Clustered indexes are more preferable than nonclustered, if you need to select by a range of values or you need to sort results set with GROUP BY or ORDER BY.
- If your application will be performing the same query over and over on the same table, consider creating a covering index on the table.
- You can use the SQL Server Profiler Create Trace Wizard with "Identify Scans of Large Tables" trace to determine which tables in your database may need indexes. This trace will show which tables are being scanned by queries instead of using an index.
- You can use sp_MSforeachtable undocumented stored procedure to rebuild all indexes in your database. Try to schedule it to execute during CPU idle time and slow production periods.
sp_MSforeachtable @command1="print '?' DBCC DBREINDEX ('?')"

T-SQL Queries

1. 2 tables

Employee	Phone
empid empname salary mgrid	empid phnumber

Select all employees who doesn't have phone?

```
SELECT empname
FROM Employee
WHERE (empid NOT IN
(SELECT DISTINCT empid
FROM phone))
```

3. Select the employee names who is having more than one phone numbers.

```
SELECT empname
FROM employee
WHERE (empid IN
(SELECT empid
FROM phone
GROUP BY empid
HAVING COUNT(empid) > 1))
```

4. Select the details of 3 max salaried employees from employee table.

```
SELECT TOP 3 empid, salary
FROM employee
ORDER BY salary DESC
```

5. Display all managers from the table. (manager id is same as emp id)

```
SELECT empname
FROM employee
WHERE (empid IN
(SELECT DISTINCT mgrid
FROM employee))
```

Write a Select statement to list the Employee Name, Manager Name under a particular manager?

```
SELECT e1.empname AS EmpName, e2.empname AS ManagerName
FROM Employee e1 INNER JOIN
Employee e2 ON e1.mgrid = e2.empid
ORDER BY e2.mgrid
```

2 tables emp and phone.

emp fields are - empid, name

Ph fields are - empid, ph (office, mobile, home). Select all employees who doesn't have any ph nos.

```
SELECT *
FROM employee LEFT OUTER JOIN
phone ON employee.empid = phone.empid
WHERE (phone.office IS NULL OR phone.office = ' ')
AND (phone.mobile IS NULL OR phone.mobile = ' ')
AND (phone.home IS NULL OR phone.home = ' ')
```

8. Find employee who is living in more than one city.

Two Tables:

Emp	City
Empid empName Salary	Empid City

9. SELECT empname, fname, lname

```
FROM employee
WHERE (empid IN
(SELECT empid
FROM city
GROUP BY empid
HAVING COUNT(empid) > 1))
```

Find all employees who is living in the same city. (table is same as above)

```
SELECT fname
FROM employee
WHERE (empid IN
(SELECT empid
FROM city a
WHERE city IN
(SELECT city
FROM city b
GROUP BY city
HAVING COUNT(city) > 1)))
```

There is a table named MovieTable with three columns - moviename, person and role. Write a query which gets the movie details where Mr. Amitabh and Mr. Vinod acted and their role is actor.

```
SELECT DISTINCT m1.moviename
FROM MovieTable m1 INNER JOIN
MovieTable m2 ON m1.moviename = m2.moviename
WHERE (m1.person = 'amitabh' AND m2.person = 'vinod' OR
m2.person = 'amitabh' AND m1.person = 'vinod') AND (m1.role =
'actor') AND (m2.role = 'actor')
ORDER BY m1.moviename
```

There are two employee tables named emp1 and emp2. Both contains same structure (salary details). But Emp2 salary details are incorrect and emp1 salary details are correct. So, write a query which corrects salary details of the table emp2

```
update a set a.sal=b.sal from emp1 a, emp2 b where a.empid=b.empid
```

13. Given a Table named "Students" which contains studentid, subjectid and marks. Where there are 10 subjects and 50 students. Write a Query to find out the Maximum marks obtained in each subject.

14. In this same tables now write a SQL Query to get the studentid also to combine with previous results.

Three tables – student , course, marks – how do go at finding name of the students who got max marks in the diff courses.

```
SELECT student.name, course.name AS coursename, marks.sid,
marks.mark
FROM marks INNER JOIN
student ON marks.sid = student.sid INNER JOIN
course ON marks.cid = course.cid
WHERE (marks.mark =
(SELECT MAX(Mark)
FROM Marks MaxMark
WHERE MaxMark.cID = Marks.cID))
```

There is a table day_temp which has three columns dayid, day and temperature. How do I write a query to get the difference of temperature among each other for seven days of a week?

```
SELECT a.dayid, a.dday, a.tempe, a.tempe - b.tempe AS Difference
FROM day_temp a INNER JOIN
day_temp b ON a.dayid = b.dayid + 1
```

OR

Select a.day, a.degree-b.degree from temperature a, temperature b
where a.id=b.id+1

There is a table which contains the names like this. a1, a2, a3, a3, a4,
a1, a1, a2 and their salaries. Write a query to get grand total salary, and
total salaries of individual employees in one query.

```
SELECT empid, SUM(salary) AS salary  
FROM employee  
GROUP BY empid WITH ROLLUP  
ORDER BY empid
```

How to know how many tables contains empno as a column in a database?

```
SELECT COUNT(*) AS Counter  
FROM syscolumns  
WHERE (name = 'empno')
```

Find duplicate rows in a table? OR I have a table with one column which has many records which are not distinct. I need to find the distinct values from that column and number of times it's repeated.

```
SELECT sid, mark, COUNT(*) AS Counter  
FROM marks  
GROUP BY sid, mark  
HAVING (COUNT(*) > 1)
```

How to delete the rows which are duplicate (don't delete both duplicate records).

```
SET ROWCOUNT 1  
DELETE yourtable  
FROM yourtable a  
WHERE (SELECT COUNT(*) FROM yourtable b WHERE b.name1 =  
a.name1 AND b.age1 = a.age1) > 1  
WHILE @@rowcount > 0  
    DELETE yourtable  
    FROM yourtable a  
    WHERE (SELECT COUNT(*) FROM yourtable b WHERE b.name1 =  
a.name1 AND b.age1 = a.age1) > 1  
SET ROWCOUNT 0
```

How to find 6th highest salary

```
SELECT TOP 1 salary  
FROM (SELECT DISTINCT TOP 6 salary  
FROM employee  
ORDER BY salary DESC) a  
ORDER BY salary
```

Find top salary among two tables

```
SELECT TOP 1 sal  
FROM (SELECT MAX(sal) AS sal  
FROM sal1  
UNION  
SELECT MAX(sal) AS sal
```

```
FROM sal2) a
ORDER BY sal DESC
```

Write a query to convert all the letters in a word to upper case
SELECT UPPER('test')

Write a query to round up the values of a number. For example even if the user enters 7.1 it should be rounded up to 8.

```
SELECT CEILING (7.1)
```

Write a SQL Query to find first day of month?

```
SELECT DATENAME(dw, DATEADD(dd, - DATEPART(dd, GETDATE()) + 1, GETDATE())) AS FirstDay
```

Datepart	Abbreviations
year	yy, yyyy
quarter	qq, q
month	mm, m
dayofyear	dy, y
day	dd, d
week	wk, ww
weekday	dw
hour	hh
minute	mi, n
second	ss, s
millisecond	ms

Table A contains column1 which is primary key and has 2 values (1, 2) and Table B contains column1 which is primary key and has 2 values (2, 3). Write a query which returns the values that are not common for the tables and the query should return one column with 2 records.

```
SELECT tbla.a
FROM tbla, tblb
WHERE tbla.a <>
(SELECT tblb.a
FROM tbla, tblb
WHERE tbla.a = tblb.a)
UNION
SELECT tblb.a
FROM tbla, tblb
WHERE tblb.a <>
(SELECT tbla.a
FROM tbla, tblb
WHERE tbla.a = tblb.a)
```

OR (better approach)

```
SELECT a
FROM tbla
WHERE a NOT IN
(SELECT a
FROM tblb)
```

```

UNION ALL
SELECT a
FROM tblb
WHERE a NOT IN
(SELECT a
FROM tbla)

```

There are 3 tables Titles, Authors and Title-Authors (check PUBS db). Write the query to get the author name and the number of books written by that author, the result should start from the author who has written the maximum number of books and end with the author who has written the minimum number of books.

```

SELECT authors.au_lname, COUNT(*) AS BooksCount
FROM authors INNER JOIN
titleauthor ON authors.au_id = titleauthor.au_id INNER JOIN
titles ON titles.title_id = titleauthor.title_id
GROUP BY authors.au_lname
ORDER BY BooksCount DESC

```

```

UPDATE emp_master
SET emp_sal =
CASE
WHEN emp_sal > 0 AND emp_sal <= 20000 THEN (emp_sal * 1.01)
WHEN emp_sal > 20000 THEN (emp_sal * 1.02)
END

```

List all products with total quantity ordered, if quantity ordered is null show it as 0.

```

SELECT name, CASE WHEN SUM(qty) IS NULL THEN 0 WHEN SUM(qty)
> 0 THEN SUM(qty) END AS tot
FROM [order] RIGHT OUTER JOIN
product ON [order].prodid = product.prodid
GROUP BY name

```

Result:

```

coke 60
mirinda 0
pepsi 10

```

30. ANY, SOME, or ALL?

ALL means greater than every value--in other words, greater than the maximum value. For example, >ALL (1, 2, 3) means greater than 3.

ANY means greater than at least one value, that is, greater than the minimum. So >ANY (1, 2, 3) means greater than 1. SOME is an SQL-92 standard equivalent for ANY.

IN & = (difference in correlated sub query)

INDEX

What is Index? It's purpose?

Indexes in databases are similar to indexes in books. In a database, an index allows the database program to find data in a table without scanning the entire table. An index in a database is a list of values in a table with the storage locations of rows in the table that contain each

value. Indexes can be created on either a single column or a combination of columns in a table and are implemented in the form of B-trees. An index contains an entry with one or more columns (the search key) from each row in a table. A B-tree is sorted on the search key, and can be searched efficiently on any leading subset of the search key. For example, an index on columns **A, B, C** can be searched efficiently on **A**, on **A, B**, and **A, B, C**.

Explain about Clustered and non clustered index? How to choose between a Clustered Index and a Non-Clustered Index?

There are clustered and nonclustered indexes. A clustered index is a special type of index that reorders the way records in the table are physically stored. Therefore table can have only one clustered index. The leaf nodes of a clustered index contain the data pages.

A nonclustered index is a special type of index in which the logical order of the index does not match the physical stored order of the rows on disk. The leaf nodes of a nonclustered index does not consist of the data pages. Instead, the leaf nodes contain index rows.

Consider using a clustered index for:

- Columns that contain a large number of distinct values.
 - Queries that return a range of values using operators such as BETWEEN, >, >=, <, and <=.
 - Columns that are accessed sequentially.
 - Queries that return large result sets.
- Non-clustered indexes have the same B-tree structure as clustered indexes, with two significant differences:
- The data rows are not sorted and stored in order based on their non-clustered keys.
 - The leaf layer of a non-clustered index does not consist of the data pages. Instead, the leaf nodes contain index rows. Each index row contains the non-clustered key value and one or more row locators that point to the data row (or rows if the index is not unique) having the key value.
 - Per table only 249 non clustered indexes.

Disadvantage of index?

Every index increases the time it takes to perform INSERTS, UPDATES and DELETES, so the number of indexes should not be very much.

Given a scenario that I have a 10 Clustered Index in a Table to all their 10 Columns. What are the advantages and disadvantages?

A: Only 1 clustered index is possible.

How can I enforce to use particular index?

You can use index hint (index=<index_name>) after the table name.

```
SELECT au_lname FROM authors (index=au_nmind)
```

What is Index Tuning?

One of the hardest tasks facing database administrators is the selection of appropriate columns for non-clustered indexes. You should consider creating non-clustered indexes on any columns that are frequently referenced in the WHERE clauses of SQL statements. Other good candidates are columns referenced by JOIN and GROUP BY operations. You may wish to also consider creating non-clustered indexes that cover

all of the columns used by certain frequently issued queries. These queries are referred to as "covered queries" and experience excellent performance gains.

Index Tuning is the process of finding appropriate column for non-clustered indexes.

SQL Server provides a wonderful facility known as the Index Tuning Wizard which greatly enhances the index selection process.

Difference between Index defrag and Index rebuild?

When you create an index in the database, the index information used by queries is stored in index pages. The sequential index pages are chained together by pointers from one page to the next. When changes are made to the data that affect the index, the information in the index can become scattered in the database. Rebuilding an index reorganizes the storage of the index data (and table data in the case of a clustered index) to remove fragmentation. This can improve disk performance by reducing the number of page reads required to obtain the requested data

DBCC INDEXDEFRAG - Defragments clustered and secondary indexes of the specified table or view.

**

What is sorting and what is the difference between sorting & clustered indexes?

The ORDER BY clause sorts query results by one or more columns up to 8,060 bytes. This will happen by the time when we retrieve data from database. Clustered indexes physically sorting data, while inserting/updating the table.

What are statistics, under what circumstances they go out of date, how do you update them?

Statistics determine the selectivity of the indexes. If an indexed column has unique values then the selectivity of that index is more, as opposed to an index with non-unique values. Query optimizer uses these indexes in determining whether to choose an index or not while executing a query.

Some situations under which you should update statistics:

- 1) If there is significant change in the key values in the index
- 2) If a large amount of data in an indexed column has been added, changed, or removed (that is, if the distribution of key values has changed), or the table has been truncated using the TRUNCATE TABLE statement and then repopulated
- 3) Database is upgraded from a previous version

What is fillfactor? What is the use of it ? What happens when we ignore it? When you should use low fill factor?

When you create a clustered index, the data in the table is stored in the data pages of the database according to the order of the values in the indexed columns. When new rows of data are inserted into the table or the values in the indexed columns are changed, Microsoft® SQL Server™ 2000 may have to reorganize the storage of the data in the table to make room for the new row and maintain the ordered storage of the data. This also applies to nonclustered indexes. When data is added

or changed, SQL Server may have to reorganize the storage of the data in the nonclustered index pages. When a new row is added to a full index page, SQL Server moves approximately half the rows to a new page to make room for the new row. This reorganization is known as a page split. Page splitting can impair performance and fragment the storage of the data in a table.

When creating an index, you can specify a fill factor to leave extra gaps and reserve a percentage of free space on each leaf level page of the index to accommodate future expansion in the storage of the table's data and reduce the potential for page splits. The fill factor value is a percentage from 0 to 100 that specifies how much to fill the data pages after the index is created. A value of 100 means the pages will be full and will take the least amount of storage space. This setting should be used only when there will be no changes to the data, for example, on a read-only table. A lower value leaves more empty space on the data pages, which reduces the need to split data pages as indexes grow but requires more storage space. This setting is more appropriate when there will be changes to the data in the table.

DATA TYPES

42. What are the data types in SQL

bigint	Binary	bit	char	cursor
datetime	Decimal	float	image	int
money	Nchar	ntext	nvarchar	real
smalldatetime	Smallint	smallmoney	text	timestamp
tinyint	Varbinary	Varchar	uniqueidentifier	

Difference between char and nvarchar / char and varchar data-type?

char[(n)] - Fixed-length non-Unicode character data with length of n bytes. n must be a value from 1 through 8,000. Storage size is n bytes. The SQL-92 synonym for char is character.

nvarchar(n) - Variable-length Unicode character data of n characters. n must be a value from 1 through 4,000. Storage size, in bytes, is two times the number of characters entered. The data entered can be 0 characters in length. The SQL-92 synonyms for nvarchar are national char varying and national character varying.

varchar[(n)] - Variable-length non-Unicode character data with length of n bytes. n must be a value from 1 through 8,000. Storage size is the actual length in bytes of the data entered, not n bytes. The data entered can be 0 characters in length. The SQL-92 synonyms for varchar are char varying or character varying.

GUID datasize?

128bit

How GUID becoming unique across machines?

To ensure uniqueness across machines, the ID of the network card is used (among others) to compute the number.

What is the difference between text and image data type?

Text and image. Use text for character data if you need to store more

than 255 characters in SQL Server 6.5, or more than 8000 in SQL Server 7.0. Use image for binary large objects (BLOBs) such as digital images. With text and image data types, the data is not stored in the row, so the limit of the page size does not apply. All that is stored in the row is a pointer to the database pages that contain the data. Individual text, ntext, and image values can be a maximum of 2-GB, which is too long to store in a single data row.

JOINS

What are joins?

Sometimes we have to select data from two or more tables to make our result complete. We have to perform a join.

How many types of Joins?

Joins can be categorized as:

Inner joins (the typical join operation, which uses some comparison operator like = or <>). These include equi-joins and natural joins.

Inner joins use a comparison operator to match rows from two tables based on the values in common columns from each table. For example, retrieving all rows where the student identification number is the same in both the **students** and **courses** tables.

Outer joins. Outer joins can be a left, a right, or full outer join.

Outer joins are specified with one of the following sets of keywords when they are specified in the FROM clause:

- LEFT JOIN or LEFT OUTER JOIN -The result set of a left outer join includes all the rows from the left table specified in the LEFT OUTER clause, not just the ones in which the joined columns match. When a row in the left table has no matching rows in the right table, the associated result set row contains null values for all select list columns coming from the right table.
- RIGHT JOIN or RIGHT OUTER JOIN - A right outer join is the reverse of a left outer join. All rows from the right table are returned. Null values are returned for the left table any time a right table row has no matching row in the left table.
- FULL JOIN or FULL OUTER JOIN - A full outer join returns all rows in both the left and right tables. Any time a row has no match in the other table, the select list columns from the other table contain null values. When there is a match between the tables, the entire result set row contains data values from the base tables.

Cross joins - Cross joins return all rows from the left table, each row from the left table is combined with all rows from the right table. Cross joins are also called **Cartesian products**. (A Cartesian join will get you a Cartesian product. A Cartesian join is when you join every row of one table to every row of another table. You can also get one by joining every row of a table to every row of itself.)

2. **What is self join?**

A table can be joined to itself in a self-join.

3. **What are the differences between UNION and JOINS?**

A join selects columns from 2 or more tables. A union selects rows.

4. **Can I improve performance by using the ANSI-style joins instead of the old-style joins?**

Code Example 1:

```
select o.name, i.name  
from sysobjects o, sysindexes i  
where o.id = i.id
```

Code Example 2:

```
select o.name, i.name  
from sysobjects o inner join sysindexes i  
on o.id = i.id
```

You will not get any performance gain by switching to the ANSI-style JOIN syntax.

Using the ANSI-JOIN syntax gives you an important advantage:

Because the join logic is cleanly separated from the filtering criteria, you can understand the query logic more quickly.

The SQL Server old-style JOIN executes the filtering conditions before executing the joins, whereas the ANSI-style JOIN reverses this procedure (join logic precedes filtering).

Perhaps the most compelling argument for switching to the ANSI-style JOIN is that Microsoft has explicitly stated that SQL Server will not support the old-style OUTER JOIN syntax indefinitely.

Another important consideration is that the ANSI-style JOIN supports query constructions that the old-style JOIN syntax does not support.

5. **What is derived table?**

pubs